

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
10 April 2003 (10.04.2003)

PCT

(10) International Publication Number
WO 03/029977 A2

(51) International Patent Classification⁷: **G06F 9/46**

(21) International Application Number: PCT/US02/31261

(22) International Filing Date: 1 October 2002 (01.10.2002)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
09/970,038 2 October 2001 (02.10.2001) US
09/970,037 2 October 2001 (02.10.2001) US
09/970,462 2 October 2001 (02.10.2001) US

(71) Applicant: CITRIX SYSTEMS, INC. [US/US]; 851 W.
Cypress Creek Road, Fort Lauderdale, FL 33309 (US).

(81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VC, VN, YU, ZA, ZM, ZW.

(84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— *without international search report and to be republished upon receipt of that report*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(72) Inventors: PEART, Franklyn; 3371 Overlook Road, Davie, FL 33328 (US). RICHTBERG, Michael; 5461 NW 86th Terrace, Coral Springs, FL 33067 (US).

(74) Agent: LANZA, John, D.; Testa, Hurwitz & Thibault, LLP, High Street Tower, 125 High Street, Boston, MA 02110 (US).



WO 03/029977 A2

(54) Title: METHODS FOR DISTRIBUTED PROGRAM EXECUTION WITH FILE-TYPE ASSOCIATION IN A CLIENT-SERVER NETWORK

(57) Abstract: In one aspect, a mapping between a file type and a program is received, a client system file is graphically presented, a selection of the graphical depiction is received, a program associated with the selected file is identified, and a request to execute the identified program is sent. In another aspect, a server system file is graphically presented, a selection of the graphical depiction is received, the selection is transmitted to a server system, a request to execute an associated program is received, and the associated program is executed. In still another aspect, a web server file is graphically presented and a request to execute a first program that includes information concerning a second program associated with the file is received. Application output data from the second program is reviewed and formatted by the first program and received application output data is presented to the user.

METHODS FOR DISTRIBUTED PROGRAM EXECUTION WITH FILE-TYPE ASSOCIATION IN A CLIENT-SERVER NETWORK

FIELD OF THE INVENTION

[0001] The present invention relates generally to distributed program execution in client-server networks. In particular, the present invention relates to methods for automatically executing a program associated with a data file when the data file and the executable program are
5 located on different computing nodes.

BACKGROUND OF THE INVENTION

[0002] In desktop computers employing command-line based operating systems, such as DISK OPERATING SYSTEM (DOS) from Microsoft Corporation of Redmond, Washington or UNIX from AT&T Corporation of New York, NY, initiating the execution of a computer
10 program required a user to enter the path and filename of the executable program at a command prompt. In response, the operating system attempted to execute the specified file. The user also entered one or more "arguments" after the path and filename that provided additional information that controlled the execution of the executable program. For example, typical arguments affected program operation and included the names of data files that the program
15 would process. If an argument specified a data file for processing — for example, a spreadsheet — whose contents were not supported by the executable program — e.g., a word processor — the result ranged from system failure to a scrambled display. This problem also existed between similar software sold by different vendors: a word processor from one vendor typically was incapable of reading the files associated with another vendor's word processor.

20 [0003] One solution to this problem involved a voluntary naming convention: filenames ended in a period and an extension that specified their contents. However, nothing ensured that different vendors would not adopt the same file extension, that the filename extension was correct, or that the contents of the file were accurately specified by the extension. Moreover, nothing ensured that a user, confronted with a file with an unfamiliar extension, would utilize the

- 2 -

correct executable program to access the contents of the file, assuming the user had the appropriate executable program to begin with.

[0004] The MACINTOSH OPERATING SYSTEM (MacOS) from Apple Computer of Cupertino, CA introduced one solution to this problem for desktop computers. Each file stored on a computer running MacOS had a file type and a creator identifier associated with it. Each executable file under MacOS had a special file type indicating that the file was executable and an associated creator identifier that matched the creator identifier of the files it was capable of editing. To edit a data file, a user selected a graphical depiction of the data file, which automatically initiated execution of an associated executable program sharing the creator identifier of the data file. If the computer lacked an executable program with a creator identifier matching that of the selected data file, the user was presented with a dialog informing her of that fact.

[0005] Moving beyond the desktop, contemporary computer networks consist of a number of computer systems, called nodes, communicating with other computer systems via communications links. Typically, some of the nodes are client nodes and other nodes are server nodes. A client node formulates and delivers queries to a server node. A user of the client node enters the queries through a user interface operating on the client node. The server node evaluates the queries and delivers responses to the client node for display on the client user interface.

[0006] Usually, the server nodes host a variety of application programs that are accessed and executed by client nodes. When a client node launches an application program, the execution of that application program occurs at either the client node or the server node, depending upon the computing model followed by the computer network. In a server-based computing model, the server node executes the application program, and only the control information for the client user interface is transmitted across the computer network for the client node for display. In a client-based computing model, the server node transmits the application program to the client node so that the client node executes the program using the resources of the client node.

[0007] Users of contemporary computer networks desire functionality similar to that provided on desktop computers by MacOS and members of the WINDOWS family of operating systems offered by Microsoft Corporation of Redmond, Washington. However, in a client-server

- 3 -

network the executable program, the data file, and the user's virtual desktop may all be present on different computers. In this case, selection of the data file will typically invoke the operating system functionality described above, attempting to initiate the execution of an executable program on the same computer as the data file. For several reasons, this is not likely the behavior desired by the user: the computer storing the data file may not have the desired application; the computer storing the data file may be shared, subjecting other users to significant performance degradation when executing executable programs, etc. Therefore, it is desirable to provide apparatus and methods that facilitate distributed program execution, whereby selecting a data file or a representation thereof on one computer initiates the execution of an associated executable program on another computer.

SUMMARY OF THE INVENTION

[0008] The present invention relates to methods for distributed program execution. In one aspect, the invention relates to a method for distributed program execution in a network including a client system and a plurality of server systems. The method includes the step of receiving a mapping specifying an association between a type of data file and an executable program for execution on a server. The method also includes the steps of presenting a graphical depiction of a data file stored on a client system and receiving a selection of the graphical depiction of the data file. The method further includes the steps of identifying an executable program associated with the type of the selected data file using the received mapping and sending a request to execute the identified executable program.

[0009] In one embodiment, the mapping includes information related to the type of data file, a program to connect with a server, an executable program for processing data, and other parameters as input for the executable program. In another embodiment, the mapping may be updated periodically or on an as-needed basis. In yet another embodiment, a file at the client system is modified to include the received mapping. The client, from the mapping, may identify the executable program and its location and may send a request to the identified server or servers to execute the program.

[0010] In another aspect, the invention relates to a method for enabling distributed program execution in a network including a client system and a plurality of server systems. The method

- 4 -

includes providing a mapping specifying an association between a type of data file and an executable program for execution on a server system. The data file is stored on the client system. The method also includes the steps of receiving a request to execute the executable program and executing the executable program on a server.

- 5 [0011] In still another aspect, the invention relates to a method for distributed program execution in a network including a client system and a plurality of server systems. The method includes the steps of presenting a graphical depiction of a data file stored on a server system and receiving a selection of the graphical depiction. The method also includes the steps of transmitting the selection of the graphical depiction of the data file to a server system and
10 receiving a request to execute an executable program associated with the selected data file. The request is transmitted to a server system. The method additionally includes executing the executable program associated with the selected data file.

- [0012] In one embodiment, the method also includes the step of receiving data associated with the selected data file. The executing program uses the received data for processing. In
15 another embodiment, a pointer to the selected data file is received, and the pointer may be utilized to retrieve data associated with the selected data file from a server system.

- [0013] In yet another aspect, the invention relates to a method for enabling distributed program execution in a network including a client system and a plurality of server systems. The method includes the step of receiving a mapping specifying an association between a type of data
20 file and an executable program for execution on a client system. The method also includes the steps of storing a data file on a server system, receiving a selection of the stored data file, and identifying an executable program associated with the type of the selected data file using the received mapping. The method additionally includes sending a request to a client system to execute the identified executable program.

- 25 [0014] In one embodiment, the data file is modified on a server system to include the received mapping. The received mapping may also be updated on a periodic basis or on an as-needed basis.

- [0015] In yet another aspect, the invention relates to a method for enabling distributed program execution in a network that includes a client system and a plurality of server system.

- 5 -

The method includes the step of receiving a mapping specifying an association between a type of data file and an executable program for execution on a client system. The method also includes the steps of storing a data file on a server system, receiving a selection of the stored data file at a client system, and transmitting the selection of the data file to a server system. The method
5 further includes the step of identifying an executable program associated with the type of the selected data file using the received mapping specifying the association. The method additionally includes sending a request to a client system to execute the identified executable program and executing the executable program at the client system.

[0016] In another aspect the invention relates to a method for enabling distributed program
10 execution in a network that includes a client system and a plurality of server systems. The method includes the step of presenting a graphical depiction of a data file that is stored on a web server. A request to execute a first executable program that includes information concerning a second executable program associated with the data file is received. The second program executes on the client system or one of a plurality of server systems. Application output data
15 from the second executable program is received and formatted by the first executable program for presentation. In other embodiments, the method also includes the step of receiving the first executable program. In still other embodiments the method includes the step of receiving a rule for determining whether an identified executable program is to be executed on the client system or one of the plurality of server systems. In still further embodiments, the method includes the
20 step of displaying the application output data using the first executable program. In some of these embodiments the method includes the step of receiving data associated with the data file.

[0017] In another aspect, the invention relates to a method for enabling distributed program execution in a network including a client system and a plurality of server systems. A graphical depiction of a data file is transmitted and a request to execute a first executable program is made.
25 The first executable program receives application output data. The request may include information concerning a second executable program for execution on either the client system or one of a plurality of server systems. In one embodiment, the method includes step of providing the first executable program or a pointer to the data file. In yet another embodiment a method includes a step of servicing requests for data associated with the data file utilizing the provided
30 pointer. In still another embodiment, the method includes a step of providing the contents of the data file.

- 6 -

[0018] In yet another aspect, the invention relates to a method for enabling distributed program execution in a network that includes a client system and a plurality of server systems. The method includes the step of transmitting, by a web system, a graphical depiction of a data file stored on a web server. The client system is provided with a request to execute a first
5 executable program for receiving application output. The request includes information concerning a second executable program for execution on the client system or one of a plurality of server systems. The client system executes the first executable program. In one embodiment the method includes the step of providing, to the client system, the first executable program. In a further embodiment the method includes a step of receiving, at the client system, a rule for
10 determining whether the second executable program is to be executed on the client system or one of the plurality of server systems.

BRIEF DESCRIPTION OF THE DRAWINGS

[0019] The invention is pointed out with particularity in the appended claims. The above and further advantages of the invention may be better understood by referring to the following
15 description and the accompanying drawings, in which:

[0020] FIG. 1A is a diagram of an embodiment of client nodes in communications with a group of server nodes via a network, wherein a program neighborhood of a client node can be determined;

[0021] FIG. 1B is an illustration of the display of an embodiment of a client node employing
20 a graphical user interface;

[0022] FIG. 2A is a block diagram illustrating an exemplary process by which one of the server nodes initiates execution of an application program for determining the program neighborhood of a client node;

[0023] FIG. 2B is a block diagram illustrating an exemplary process by which a client node
25 initiates execution of an application program for determining the program neighborhood of that client node;

- 7 -

[0024] FIG. 2C is a block diagram illustrating an exemplary process by which a client node uses a web browser application to determine its program neighborhood;

[0025] FIGS. 3A, 3B, and 3C are block diagrams illustrating exemplary processes by which a client node launches an application program from a Program Neighborhood window displayed at
5 that client node;

[0026] FIG. 3D is a block diagram illustrating an exemplary process by which a client node launches an application program from a Program Neighborhood web page displayed at that client node;

[0027] FIG. 4 is a block diagram of a client-based computing embodiment in which a client
10 node having an installed program neighborhood application is in communication with one of a plurality of server nodes;

[0028] FIG. 5 is a block diagram of a server-based computing embodiment in which a client node is in communication with a server node having an installed program neighborhood application;

[0029] FIG. 6A is a screenshot of an exemplary display on the display screen of a client node
15 after the program neighborhood application program is executed;

[0030] FIG. 6B is a screenshot of another exemplary display on the display screen of a client node after the program neighborhood application program of the invention is executed;

[0031] FIG. 7 is a flow chart representation of an embodiment of a process by which a client
20 node is informed as to the availability for use of application programs on the application servers;

[0032] FIG. 8A is a flow chart representation of an embodiment of a process by which a user executes an application program on a server node by selecting a data file hosted by his or her client node;

[0033] FIG. 8B is a flow chart representation of the role of the server node in the
25 embodiment of FIG. 8A;

- 8 -

[0034] FIG. 9A is a flow chart representation of an embodiment of a process by which a user executes an application program on his or her client node by selecting a data file hosted by a server node;

5 [0035] FIG. 9B is a flow chart representation of the role of the server node in the embodiment of FIG. 9A;

[0036] FIG. 10A is a flow chart representation of an embodiment of a process by which a user executes an application hosted by a server node by selecting a data file hosted by a web server; and

10 [0037] FIG. 10B is a flow chart representation of the role of the web server in the embodiment of FIG. 10A.

[0038] In the drawings, like reference characters generally refer to the same parts throughout the different views. The drawings are not necessarily to scale, emphasis instead being placed upon illustrating the principles of the invention.

DETAILED DESCRIPTION OF THE INVENTION

15 [0039] FIG. 1 shows a first computing system (client node) 10 and a second computing system (client node) 20 in communication with computing systems (application servers) 30, 32, 34, and 36 over a network 40. The network 40 can be a local-area network (LAN) or a wide area network (WAN) such as the Internet or the World Wide Web. Users of the client nodes 10, 20 can be connected to the network 40 through a variety of connections including standard
20 telephone lines, LAN or WAN links (e.g., T1, T3, 56kb, X.25), broadband connections (ISDN, Frame Relay, ATM), and wireless connections. The connections can be established using a variety of communication protocols (e.g., TCP/IP, IPX, SPX, NetBIOS, Ethernet, RS232, and direct asynchronous connections).

[0040] In one embodiment, the client node 10 is a personal computer (e.g. using a
25 microprocessor from the x86, 680x0, PowerPC, PA-RISC, or MIPS processor families), smart or dumb terminal, network computer, wireless device, information appliance, workstation, minicomputer, mainframe computer or other computing device that has a graphical user interface

- 9 -

and sufficient persistent storage to execute application programs downloaded from application servers 30, 32, 34 over the network 40. Operating systems supported by the client node 10 can include any member of the WINDOWS family of operating systems, MacOS, JavaOS, and various varieties of Unix (e.g., Solaris, SunOS, Linux, HP-UX, A/IX, and BSD-based distributions). The client node 10 can include a display 12, a keyboard 14, a memory 16 for storing downloaded application programs, a processor 17, and a mouse 18. The memory 16 can provide persistent or volatile storage.

[0041] The processor 17 can execute the application programs locally on the client node 10 and display a resulting windows-based desktop on the display screen 12. Such local processing on the client node 10 is according to the above-described client-based computing model. A typical windows-based desktop appearing on one embodiment of the client node 10 is depicted in FIG. 1B. The resources available to the client node 20 are presented graphically as one or more icons 100. Memory 16 is presented as a hard disk icon 104. The downloaded application programs and data files stored on memory 16 are presented as individual file icons. For example, memory 16 contains a downloaded application program implementing a word processor, illustrated as icon 108, and a data file containing a document adapted for processing by the word processor program, illustrated as icon 112.

[0042] The user typically manipulates the desktop using keyboard 14, mouse 18, or some other input device such as a haptic interface (not shown). Using the input device, the user invokes the execution of the word processor program by selecting icon 108. Using the word processor program, the user manipulates the contents of the data file containing the document represented by icon 112. However, it is more likely that the user will eschew locating and activating the word processor program in favor of selecting the data file itself. As described above, most modern operating systems with graphical user interfaces will examine the selected data file, retrieve one or more parameters associated with the selected data file, identify an executable application program on the desktop computer associated with the selected data file, begin execution of the associated application program, and then provide the contents of the selected data file to the associated application program for processing.

[0043] In other embodiments, the client node 20 is any terminal (windows or non-windows based), or thin-client device operating according to a server-based computing model. In a server-

- 10 -

based computing model, the execution of application programs occurs entirely on the application servers 30, 32, 34, and the user interface, keystrokes, and mouse movements are transmitted over the network 40 to the client node 20. The user interface can be text driven (e.g., DOS) or graphically driven (e.g., WINDOWS). Platforms supported by the client node 20 include DOS and WINDOWS CE for windows-based terminals. The client node 20 includes a display screen 22, a keyboard 24, a mouse 28, a processor (not shown), and persistent storage (not shown).

[0044] The application servers 30, 32, 34, and 36 are any computing device that controls access to other portions of the network (e.g., workstations, printers). Four application servers are depicted for convenience of discussion; it is to be understood that more or fewer application servers are connected to the network 40 in various embodiments of the present invention. The servers 30, 32, 34, and 36 operate according to a client-based computing model or a server-based computing model as described above.

[0045] Each application server 30, 32, 34, and 36 hosts one or more application programs that for access by the client nodes 10 and 20. Applications made available to client nodes for use are referred to as published applications. Examples of such applications include word processing programs such as WORD and spreadsheet programs such as EXCEL, both offered by Microsoft Corporation of Redmond, Washington, financial reporting programs, customer registration programs, technical support programs, customer database programs, or application set managers.

[0046] In one embodiment, the servers 30, 32, and 34 belong to the same domain 38. In the network 40, a domain is a sub-network comprising a group of application servers and client nodes under the control of one security database. In one embodiment, a domain includes one or more "server farms," i.e., a group of servers that are linked together to act as a single server system to provide centralized administration. Conversely, a server farm can include one or more domains. For servers of two different domains to belong to the same server farm, a trust relationship may need to exist between the domains specifying an association between the different domains that allows a user to access the resources associated with each domain using a single authentication credential.

[0047] In one embodiment, the application server 36 is in a different domain than the domain 38. In another embodiment, the application server 36 is in the same domain as servers 30, 32, and 34. In either embodiment, application servers 30, 32, and 34 can belong to one server farm,

- 11 -

while the server 36 belongs to another server farm, or all of the application servers 30, 32, 34, and 36 can belong to the same server farm. When a new server is connected to the network 40, the new server either joins an existing server farm or starts a new server farm.

[0048] In one embodiment, the network 40 includes a master server node for performing load-level balancing among the application servers 30, 32, 34, and 36. In another embodiment, the master server node is one of the application servers 30, 32, 34, or 36. The master server node includes a list of server addresses and load information corresponding to each of the other application servers. Using the list of available servers and their corresponding load levels, the master server node directs the client node to a particular server node on which to execute an application based on the list of available servers and corresponding load levels. Alternatively, the application servers 30, 32, 34, 36 collaborate among themselves in a peer-to-peer fashion to exchange administration information such as load levels, allowing any server 30, 32, 34, 36 to respond to a request made by a client node 10, 20.

PROGRAM NEIGHBORHOOD

[0049] According to the principles of the invention, a user of either client node 10, 20 learns of the availability of application programs hosted by the application servers 30, 32, 34, and 36 in the network 40 without requiring the user to know where to find such applications or to enter technical information necessary to link to such applications. These available application programs comprise the "program neighborhood" of the user. A system for determining a program neighborhood for a client node includes an application program (hereafter referred to as the "Program Neighborhood" application), memory for storing components of the application program, and a processor for executing the application program.

[0050] The Program Neighborhood application is installed in memory of the client node 10 and/or on the application servers 30, 32, 34, and 36 as described below. The Program Neighborhood application is a collection of services, application programming interfaces (APIs), and user interface (UI) programs that disclose to users of the client nodes 10, 20 those application programs hosted by the application servers that each client node is authorized to use (e.g., execute).

- 12 -

[0051] An application server operating according to the Program Neighborhood application collects application-related information from each of the application servers in a server farm. The application-related information for each hosted application includes but is not limited to the address of the server hosting that application, the application name, the users or groups of users who are authorized to use that application, the data types supported by the application, any file types associated with the supported data types, any file types associated with the application, conditional rules specifying the execution of the application either at a server node or at the client node, and the minimum capabilities required of the client node before establishing a connection to run the application. For example, the application may stream video data, and therefore a required minimum capability is that the client node supports video data. Other exemplary requirements are that the client node support audio data or encrypted data. The application-related information can be stored in a database as described below.

[0052] When a client node connects with the network 40, the user of the client node provides user credentials. User credentials typically include the user's name, password, and the domain name for which the user is authorized. The user's credentials are obtained from smart cards, time-based tokens, social security numbers, user passwords, personal identification (PIN) numbers, digital certificates based on symmetric key or elliptic curve cryptography, biometric characteristics of the user, or any other means by which the identification of the user of the client node is obtained and submitted for authentication.

[0053] The server responding to the client node authenticates the user based on the user credentials. In one embodiment, the user credentials are stored wherever the Program Neighborhood application is executing. For example, in one embodiment where the client node executes the Program Neighborhood application, the user credentials are also stored at the client node. In another embodiment where an application server is executing the Program Neighborhood application, the user credentials are stored at that server.

[0054] From the user credentials and the application-related information, the server also determines which application programs hosted by the application servers are available for use by the user of the client node. The server transmits information representing the available application programs to the client node. This process eliminates the need for a user of the client

- 13 -

node to set-up application connections. Also, an administrator of the server can control access to applications among the various client node users.

[0055] The user authentication performed by the server can suffice to authorize the use of each hosted application program presented to the client node, although such applications may reside at another server. Accordingly, when the client node initiates execution of one of the hosted applications, additional input of user credentials by the user may be unnecessary to authenticate use of that application. Thus, a single entry of the user credentials serves to determine the available applications and to authorize the launching of such applications without an additional, manual log-on authentication process by the client user.

10 [0056] Either a client node 10, 20 or an application server can launch the Program Neighborhood application as described in connection with Figs. 2A-2C. The results are displayed on the display screen 12, 22 of the client node 10, 20. In a graphical environment, the results can be displayed in a Program Neighborhood graphical window and each authorized application program can be represented by a graphical icon in that window.

15 [0057] One embodiment of the Program Neighborhood application filters those application programs that the client node 10, 20 is unauthorized to use and displays only authorized (i.e., available) programs. In other embodiments, the Program Neighborhood application displays authorized and unauthorized applications. When unauthorized applications are not filtered from the display, a notice can be provided indicating that such applications are unavailable.

20 Alternatively, the Program Neighborhood application can report all applications hosted by the application servers 30, 32, 34, 36 to the user of a client node, without identifying which applications the client node 10, 20 is authorized or unauthorized to execute. Authorization is subsequently determined when the client node 10, 20 attempts to run one of those applications.

[0058] FIG. 2A shows an exemplary process by which a server launches the Program Neighborhood (PN) application and presents results of the PN application to the client node 10.

25 The server launches the PN application in response to a request 42 by the client node 10 for a particular application program. The request passes to the master server node, in this example server 30. The master server node 30, taking load-balancing and application availability into account, indicates (arrow 43) to the client node 10 that the sought-after application is available

30 on server 32. The client node 10 and server 32 establish a connection (arrows 45 and 46). By

- 14 -

this connection, the server 32 can transfer the executable code of the particular application to the client node 10, when the client node 10 and server 32 are operating according to the client-based computing model. Alternatively, the server 32 executes the particular application and transfers the graphical user interface to the client node 10, when the client node 10 and server 32 operate according to the server-based computing model. In addition, either the master server node 30 or the server 32 can execute the Program Neighborhood application 41 and push the results (arrows 43 or 46) back to the client node 10 so that when the client node 10 requests the Program Neighborhood application, the program neighborhood is already available at the client node 10.

[0059] FIG. 2B shows another exemplary process by which the client node 10 initiates execution of the Program Neighborhood application and a server presents the results of the PN application to the client node 10. The client node 10 launches the Program Neighborhood application (e.g., by selecting the Program Neighborhood icon 47 representing the application). The request 50 for the Program Neighborhood application is directed to the master server node, in this example server 30. The master server node 30 can execute the Program Neighborhood application if the application is on the master server node 30 and return the results to the client node 10. Alternatively, the master server node 30 can indicate (arrow 51) to the client node 10 that the Program Neighborhood application 41 is available on another server, in this example server 32. The client node 10 and server 32 establish a connection (arrows 53 and 54) by which the client node 10 requests execution of the Program Neighborhood application 41. The server 32 can execute the application 41 and transfer the results (i.e., the graphical user interface) to the client node 10.

[0060] FIG. 2C shows another exemplary process by which a client node 10 initiates execution of the Program Neighborhood application, in this example via the World Wide Web. A client node 20 executes a web browser application 80, such as INTERNET EXPLORER, manufactured by Microsoft Corporation of Redmond, Washington. The client node 20, via the web browser 80, transmits a request 82 to access a Uniform Resource Locator (URL) address corresponding to an HTML page residing on server 30. In some embodiments the first HTML page returned 84 to the client node 20 by the server 30 is an authentication page that seeks to identify the client node 20.

- 15 -

[0061] The authentication page allows the client node 20 to transmit user credentials via the web browser 80 to the server 30 for authentication. Transmitted user credentials are verified either by the server 30 or by another server in the farm. This allows a security domain to be projected onto the server 30. For example, if the server 30 runs the WINDOWS NT operating system and the authenticating server runs the UNIX operating system, the UNIX security domain may be said to have been projected onto the server 30. User credentials may be transmitted "in the clear," or they may be encrypted. For example, user credentials may be transmitted via a Secure Socket Layer (SSL) connection which encrypts data using the RC3 algorithm, manufactured by RSA Data Security, Inc. of San Mateo, California.

10 [0062] The server 30 may verify the user credentials received from the client node 20. Alternatively, the server 30 may pass the user credentials to another server for authentication. In this embodiment, the authenticating server may be in a different domain from the server 30. Authenticated user credentials of the client node 20 may be stored at the client node 20 in a per-session cookie, in fields that are not displayed by the web browser 80, or in any other manner
15 common in maintenance of web pages. In some embodiments, a server farm with which the server 30 is associated may allow guest users, i.e., users that do not have assigned user credentials, to access applications hosted by servers in the farm. In these embodiments, the authentication page may provide a mechanism for allowing a client node 20 to identify that it is a guest user, such as a button or menu selection. In other embodiments, the server 30 may omit the
20 authentication page entirely.

[0063] Still referring to FIG. 2C, once the client node 20 is authenticated by the server 30, the server prepares and transmits to the client node 20 an HTML page 88 that includes a Program Neighborhood window 58 in which appears graphical icons 57, 57' representing application programs to which the client node 20 has access. A user of client node 20 invokes execution of
25 an application represented by icon 57 by selecting icon 57.

[0064] FIG. 3A shows an exemplary process of communication among the client node 10, the master server node, in this example server 30, and the server 32. The client node 10 has an active connection 72 with the server 32. The client node 10 and server 32 can use the active connection 72 to exchange information regarding the execution of a first application program.

- 16 -

The user credentials of the client node 10 are stored at the client node. Such storage of the user credentials can be in cache memory or persistent storage.

[0065] In this embodiment, the Program Neighborhood application runs on the client node 10. The client node display has a Program Neighborhood window 58 which presents a graphical icon 57 representing a second application program. A user of the client node 10 can launch the second application program by selecting the icon 57 with a mouse or other input device. The request passes to the master server node 30 via a connection 59. The master server node 30 indicates to the client node 10 via the connection 59 that the sought-after application is available on server 32. The client node 10 signals the server 32 to establish a second connection 70. The server 32 requests the user credentials from the client node 10 to authenticate access to the second application program. Upon a successful authentication, the client node 10 and server 32 establish the second connection 70 and exchange information regarding the execution of the second application program. Accordingly, the client node 10 and the server 32 communicate with each other over multiple connections.

[0066] FIG. 3B shows an exemplary process of communication among the client node 20, the master server node, in this example server 30, and servers 32, 34, and 36. The client node 20 has an active connection 73 with the server 32. The client node 20 and server 32 can use the active connection 73 to exchange information regarding the execution of a first application program. The user credentials of the client node 20 are stored at the server 32 in cache memory or in persistent storage.

[0067] In this embodiment, the Program Neighborhood application runs on the server 32. The server 32 includes software providing a server-based client engine 62, enabling the server 32 to operate in the capacity of the client node 20. The client node 20 display has a Program Neighborhood window 58 in which appears graphical icons 57, 57' representing a second application program and a third application program, respectively. A user of the client node 20 can launch the second application program by selecting the icon 57. The request to launch the second application program passes to the server 32 via active connection 73, and the server 32 forwards the request to the master server node 30 (arrow 65).

[0068] The master server node 30 indicates (arrow 65) to the server 32 that the sought-after application is available on server 34. The server 32 contacts the server 34 to establish a

- 17 -

connection 66. To authenticate access to the application, the server 34 obtains the user credentials of the client node 20 from the server 32. The server 32 and server 34 establish the connection (arrow 66) by which the server 32 requests execution of the second application and the server 34 returns the graphical user interface results to the server 32. The server 32 forwards the graphical user interface results to the client node 20, where the results are displayed. Accordingly, the information exchanged between the client node 20 and the server 34 passes through the server 32.

[0069] Similarly, the client node 20 can launch the third application program by selecting the icon 57'. The request to launch the third application program passes to the server 32. The server 32 forwards the request to the master server node 30, which considers load-balancing and application program availability to determine which server can handle the request. In this example, the master server node indicates that server 36 can run the third application program.

[0070] The server 32 and server 36 establish a connection (arrow 74) by which the server 32 requests execution of the third application program, and the server 36 returns the results to the server 32. To permit execution of the third application program, the server 36 can authenticate the user credentials of the client node 20, which are obtained from the server 32. The server 32 forwards the results to the client node 20 where the results are displayed. Accordingly, the results of executing the third application program pass between the client node 20 and the server 36 through the server 32.

[0071] From this illustration it should be understood that client node 20 can run multiple application programs through one connection with the server 32, while the server 32 maintains multiple connections (in this example, one connection with server 34 and a second connection with server 36). Also, the server 32 merges the information received from the server 34 with the information received from the server 36 into one data stream for transmission to the client node 20.

[0072] FIG. 3C shows an exemplary process of communication among the client node 20, the master server node, in this example server 30, and servers 32 and 34. The client node 20 has an active connection 76 with the server 32. The client node 20 and server 32 can use the active connection 76 to exchange information regarding the execution of a first application program. The client node 20 can store the user credentials in cache memory or in persistent storage.

- 18 -

[0073] In this embodiment, the Program Neighborhood application runs on the server 32. The client node 20 display has a Program Neighborhood window 58 in which appears a graphical icon 57 representing a second application program. A user of the client node 20 can launch the second application program by selecting the icon 57. The request to launch the second application program passes to the server 32. The server 32 responds (i.e., "calls back") to the client node 20 by returning application-related information such as the name of the application and capabilities needed by the client node 20 for the second application to run.

[0074] With the information provided by the server 32, the client node 20 then communicates with the master server node 30 via connection 77 to determine the server for executing the second application program. In this example, that server is server 34. The client node 20 then establishes a connection 78 to the server 34. Server 34 requests the user credentials from the client node 20 to authenticate the user of the client node 20. The second application program executes on the server 34, and the server 34 returns the graphical user interface to the client node 20 via the established connection 78. Accordingly, the client node 20 can have multiple active connections between the multiple servers.

[0075] FIG. 3D shows an exemplary process of communicating between the client node 20, a server 30 that in this example acts as a web server, and server 32. The client node 20 authenticates itself to the server 30 as described above in connection with FIG. 2C. In one embodiment, the server 30 accesses an output display template 90, such as an SGML, HTML or XML file, to use as a base for constructing the Program Neighborhood window to transmit to the client node 20. The template may be stored in volatile or persistent memory associated with the server 30 or it may be stored in mass memory 92, such as a disk drive or optical device, as shown in FIG. 3D.

[0076] In this embodiment, the template 90 is a standard SGML, HTML, or XML document containing Program Neighborhood-specific tags that are replaced with dynamic information. The tags indicate to the server 30 where in the output display to insert information corresponding to available applications, such as icon images. In one particular embodiment, the Program Neighborhood-specific tags are embedded within comments inside the file, allowing the file to remain compatible with standard interpreters. In another embodiment, the Program

- 19 -

Neighborhood-specific tags are extensions of the markup language used as the base for the template.

[0077] Examples of HTML tags that may be used in a template in accordance with the present invention are set forth below in Table 1:

5

TABLE 1

| Tag | Description |
|--------------------------|---|
| ControlField field value | This tag is used to set the value of data that either persists between Program Neighborhood web pages, are set by the user, or are used to help in cross page navigation, such as user name, domain, password, template, and application. |
| DrawProgramNeighborhood | His tag is used to draw a Program Neighborhood display at this location in an output display. |
| AppName | This tag is replaced by the name of the published application in the current context. |
| WindowType | This tag is replaced by the window type of the published application in the current context. |
| WindowHeight | This tag is replaced by the window height of the published application in the current context. |
| WindowWidth | This tag is replaced by the window width of the published application in the current context. |
| WindowScale | This tag is replaced by the window scale of the published application in the current context. |
| WindowColors | This tag is replaced by the color depth of the published application in the current context. |
| SoundType | This tag is replaced by the sound setting of the published application in the current context. |
| VideoType | This tag is replaced by the video setting of the published application in the current context. |
| EncryptionLevel | This tag is replaced by the encryption level of the published application in the current context. |
| Icon | This tag is replaced by the icon of the published application in the current context. |

[0078] Other tags can be provided to set control fields and to provide conditional processing relating to the Program Neighborhood application.

[0079] In one embodiment, the template is constructed dynamically using, for example, COLD FUSION by Allaire Corp. of Cambridge, Massachusetts or ACTIVE SERVER PAGES by Microsoft Corporation of Redmond, Washington. Alternatively, the template may be static. The Program Neighborhood application parses the template, replacing Program Neighborhood-

10

- 20 -

specific tags as noted above. Tags that are not Program Neighborhood-specific are left in the file to be parsed by the browser program 80 executing on the client 20.

5 [0080] In one embodiment, a template parser object is provided that accepts an HTML template as input, interprets Program Neighborhood-specific tags present in the template, and provides the original template with all Program Neighborhood tags replaced with appropriate text. The template parser object can be passed a cookie, a URL query string, or a control field from a web server interface to provide the information with which Program Neighborhood-specific tags should be replaced.

10 [0081] In another embodiment, the Program Neighborhood application allows scripts to access to information via an application programming interface. Scripts may be written in, for example, VBScript or JScript. In this embodiment, the scripting language is used to dynamically generate an output display using information returned by the application in response to queries posed by the script. Once the output display is generated, it is transmitted to client node 20 for display by the browser program 80.

15 [0082] A user of the client node 20 can launch an application by selecting an icon 57, 57' displayed in the Program Neighborhood web page. In some embodiments, each icon 57, 57' is an encoded URL that specifies: the location of the application (i.e., on which servers it is hosted or, alternatively, the address of a master server); a launch command associated with the application; and a template identifying how the output of the application should be displayed
20 (i.e., in a window "embedded" in the browser or in a separate window). In some embodiments, the URL includes a file, or a reference to a file, that contains the information necessary for the client to create a connection to the server hosting the application. This file may be created by the Program Neighborhood application dynamically. The client node 20 establishes a connection (arrow 94) with the server identified as hosting the requested application (in this example, server
25 32) and exchanges information regarding execution of the desired application. In some embodiments, the connection 94 is made using the Independent Computing Architecture (ICA) protocol, manufactured by Citrix Systems, Inc. of Fort Lauderdale, Florida. Thus, the client node 20 may display application output in a window separate from the web browser 60, or it may "embed" application output within the web browser.

- 21 -

[0083] FIG. 4 illustrates an exemplary arrangement of program components for a client-based implementation of the Program Neighborhood application. A client-based implementation of Program Neighborhood application can be used in a network using either the server-based computing model in which the servers execute the Program Neighborhood application or the client-based computing model in which the client node 10 executes the Program Neighborhood application locally. The Program Neighborhood application includes a Program Neighborhood Service (PNSVC) component 44, an Application Database component 48, a Program Neighborhood Application Program Interface (PNAPI) component 52, a Program Neighborhood User Interface component 56, and a local cache 60.

10 [0084] The application server 30, for example, includes the service component (PNSVC) 44 and the application database 48. The client node 10, which is a representative example of a client node that can support a client-based implementation of the Program Neighborhood application, includes the application program interface PNAPI 52, the user interface user interface component 56, and the local cache 60 components. The PNAPI 52 communicates with
15 the user interface component 56 and the local cache 60. The PNSVC 44 communicates with the application database 48 and with the PNAPI 52 on the client node 10 via communications link 62.

[0085] The communications link 62 can be established by, for example, using the ICA protocol. ICA is a general-purpose presentation services protocol designed to run over industry
20 standard network protocols, such as TCP/IP, IPX/SPX, NetBEUI, using industry-standard transport protocols, such as ISDN, frame relay, and asynchronous transfer mode (ATM). The ICA protocol provides for virtual channels, which are session-oriented transmission connections that can be used by application-layer code to issue commands for exchanging data. The virtual channel commands are designed to be closely integrated with the functions of client nodes. One
25 type of virtual channel connection supported by the ICA protocol is a Program Neighborhood virtual channel.

[0086] The Program Neighborhood virtual channel protocol can include four groups of commands:

(1) initialization-related commands;

- 22 -

(2) single authentication related commands that can be supported by each client node wanting a copy of the user credentials;

(3) application data related commands for implementing the Program Neighborhood user interface; and

5 (4) application launch callback-related commands for running the user interface on an application server.

APPLICATION DATABASE

[0087] The application database 48 is a cache of the authorized user and group information for all the public (i.e., published) applications in a server farm or in a group of trusted domains.

10 Each server in a server farm can maintain its own application-related information in persistent storage and build up the database 48 in volatile storage. In another embodiment, all collected application-related information in the database 48 can be stored in persistent storage and made accessible to each other server in the server farm. The database 48 can be implemented in a proprietary format (e.g., as a linked list in memory) or using Novell's Directory Services (NDS)

15 or any directory service adhering to the X.500 standard defined by the International Telecommunication Union (ITU) for distributed electronic directories.

[0088] The application database 48 includes a list of application servers. Each server in the list has an associated set of applications. Associated with each application is application-related information that can include the application name, a list of servers, and client users that are

20 authorized to use that application. A simplified example of the application-related information maintained in the database is illustrated by the following Table 2. Users A and B are users of the client nodes 10, 20, "n/a" indicates that the application is hosted, but is not available to client node users, and "-" indicates that the application is not hosted.

TABLE 2

| | Applications | | | |
|-------------|--------------|-------------------|----------------|------------------|
| Server Name | Spreadsheet | Customer Database | Word Processor | Calculator |
| Server 30 | User A | User B | n/a | - |
| Server 32 | User B | n/a | User A | - |
| Server 34 | - | - | - | User A User B |

- 23 -

[0089] Table 2 shows a list of servers 30, 32, 34, the applications hosted by the servers, (Spreadsheet, Customer Database, Word Processor, and Calculator), and those users who are authorized to use the applications. For example, server 30 hosts the Spreadsheet program, the Customer Database and the Word Processor. User A is authorized to use the Spreadsheet, User B is authorized to use the Customer Database, and no users are authorized to use the Word Processor. It is to be understood that other techniques can be used to indicate who is authorized to use a particular application. For example, the user information stored in the database can be used to indicate those users who are unauthorized to use a particular application rather than those who are authorized.

[0090] To obtain the information that is stored in the database 48, the server 30 obtains the application-related information from each other server in the server farm regarding the applications on those servers, including control information that indicates which client users and servers are permitted to access each particular application. The application-related information maintained in the database may or may not persist across re-boots of the server 30.

[0091] The application database 48 can be a central database that is stored at the application servers 30 and is accessible to all of the servers in the server farm. Accordingly, the application-related information can be available for use by other servers such as those servers that perform published application authentication during session log-on and application launching. In another embodiment, the application database 48 can be maintained at each of the application servers based upon the information that each server obtains from communications with each other server in the server farm.

PROGRAM NEIGHBORHOOD SERVICE PROGRAM (PNSVC)

[0092] Each server 30, 32, 34 and 36 having the Program Neighborhood application installed thereon executes the PNSVC software 44. The PNSVC software 44, operating on each server 30, 32, 34 and 36 establishes a communication link (e.g., a named pipe) with each other server. The servers 30, 32, 34 and 36 exchange the application-related information on the named pipes. In another embodiment, the PNSVC software 44 collects the application-related information from the other servers in the server farm through remote registry calls (e.g., the service

- 24 -

component 44 transmits a datagram to other servers in the plurality requesting the application-related information corresponding to the application programs hosted by those servers). The PNSVC 44 software also maintains the relationships of groups and users to published applications in the application database 48 and accesses the information when authenticating a client user. An administrator of the server 30 can use a user interface to configure the PNSVC 44.

[0093] Other functions of the PNSVC software 44 include implementing the services and functions requested by the PNAPI 52 and communicating with the PNAPI 52 on the client node 10 using a Program Neighborhood virtual device driver (VDPN). The VDPN operates according to the Program Neighborhood virtual channel protocol described above for establishing and maintaining an ICA connection.

PROGRAM NEIGHBORHOOD APPLICATION PROGRAM INTERFACE (PNAPI)

[0094] The PNAPI 52 is a set of software functions or services that are used by the Program Neighborhood application to perform various operations (e.g., open windows on a display screen, open files, and display message boxes). The PNAPI 52 provides a generic mechanism for launching application objects (e.g., icons) produced by running the Program Neighborhood application and application objects in a legacy (i.e., predecessor or existing for some time) client user interface. When the client node 10 launches an available application, the launch mechanism can launch the application on the server 30, if necessary (e.g., when the client node 10 does not have the resources to locally execute the application).

[0095] The PNAPI 52 provides all published application information to the user interface component 56 for display on the screen 12 of the client node 10. The PNAPI 52 also manages server farm log-ons in a local database of logon credentials (e.g., passwords) for users of the client node 10 to support the single authentication feature. Credentials may or may not be persistent across bootings (power-off and on cycles) of the client node 10.

[0096] The PNAPI 52 provides automatic and manual management for Program Neighborhood application objects stored in the local cache 60. The local cache 60 can either be refreshed manually by the user of the client node 10, or at a user-definable refresh rate, or by the server at any time during a connection. In a Windows implementation, the PNAPI 52 can build

- 25 -

remote application file associations and manage the "Start" menu and desktop icons for application object shortcuts.

PROGRAM NEIGHBORHOOD USER INTERFACE

[0097] The user interface module 56 interfaces the PNAPI 52 and can be a functional
5 superset of an existing client-user interface (e.g., Remote Application Manager). The user
interface module 56 accesses the information stored in the local cache 60 through the PNAPI 52
and visually presents that information to the user on the display screen 12 of the client node 10.
The displayed information is a mixture of information generated by a user of the client node 10
and information obtained by the Program Neighborhood application. The user interface module
10 56 can also show the user all applications that the user is currently running and all active and
disconnected sessions.

[0098] In a graphical embodiment, the user interface module 56 can present a variety of
graphical components, such as windows and pull-down menus, to be displayed on the display
screen 12. A display of a combination of such graphical user interface components is generally
15 referred to as a "desktop." A desktop produced by the user interface module 56 can include a
Program Neighborhood window displaying the neighborhood of application programs available
to the user of the client node 10 for use. These application programs are a filtered combination
of the published applications hosted by a server farm on the network. The user interface module
56 can generate a Program Neighborhood window for each server farm or merge the applications
20 from different server farms under a single Program Neighborhood window.

[0099] At a top level, the Program Neighborhood window includes a folder for each server
farm. Clicking on one of the folders with the mouse 18 produces a window containing a
representation (e.g., an icon) of each hosted application available to the user, e.g., see Figs. 6A
and 6B. The Program Neighborhood window becomes the focal point for launching published
25 applications, and the user interface module 56 can be used to launch applications through the
PNAPI 52. For example, the user of the client node 10 can use the mouse 18 to select one of the
displayed icons and launch the associated application.

[00100] A feature of a client-based implementation is that the user can browse the objects
displayed in the Program Neighborhood window although the client node is offline, that is, the

- 26 -

ICA connection 62 is inactive. Also, a user of the client node 10 can drag application objects and folders out of the Program Neighborhood window and into other graphical components (e.g., other windows, folders, etc.) of the desktop.

[00101] FIG. 5 shows an exemplary arrangement of program components for a server-based implementation of the Program Neighborhood application. The components include a Service (PNSVC) component 44', an Application Database component 48', an Application Program Interface (PNAPI) component 52', a User Interface component 56' and a local cache 60'. Each software component 44', 48', 52', 56', and 60' is installed on the application server 30'. The software components for the server-based implementation correspond to the software components for the client-based implementation of FIG. 4. The functionality of each server-based software component is similar to the client-based counterpart, with differences or added capabilities described below. The PNSVC 44' communicates with the application database 48' and with the PNAPI 52' using local procedure calls. The PNAPI 52' also communicates with the user interface module 56' and the local cache 60'.

[00102] Similar to that described in FIG. 4 for the client node 10, the client node 20 logs on to the network 40, the server 30' develops and maintains a database containing the application related information collected from the other servers 32, 34 in the server farm, and a communication link is established between the server 30' and the client node 20. The application server 30' is in communication with the client node 20 via an ICA channel connection 62'. The channel connection 62' can be established by an ICA virtual channel protocol (e.g., Thinwire). The Thinwire protocol can be used to transmit presentation commands from Windows-based applications running on the application server 30' to the client node 20. To a user of the client node 20, the applications appear to be running on the client node 20. The client node 20 can include a Remote Application Manager application program 64 that communicates with the application server 30' via the ICA channel connection 62'.

[00103] To run the Program Neighborhood application in a server-based implementation, the user of the client node 20 connects to an initial desktop (at the server 30') and launches the Program Neighborhood application from within that desktop environment. The connection to the initial desktop can occur automatically, e.g., via a logon script of the client node 20, via an entry

- 27 -

in the StartUp group in Windows 95, or by another centrally managed server specific mechanism. All remote application management and launching is accomplished through this initial desktop.

[00104] Similar to that described in FIG. 4 for the server 30, the server 30' uses the user credentials to determine those application programs that are authorized for use by the user of the client node 20. A Program Neighborhood graphical window is returned to the client node 20 and displayed on the client screen 22. This window can contain icons representing the available and, possibly, the unavailable application programs that are in the program neighborhood of the client node 20.

[00105] The user of the client node 20 can select and launch one of the application programs displayed in the Program Neighborhood window. When launching an application, the Program Neighborhood application can execute the application on the same server 30', where applicable, taking into account load balancing requirements among servers and the availability of the application on that server 30'. The PNAPI 52' can include a launch mechanism for launching a remote application locally on the server 30' when the server 30' is nominated to launch the application. When a different server is needed to run the application, the Program Neighborhood application can launch the application via the server 30' (i.e., server-based client) using the windows to present the application on the desktop of the client node 20 as described above in FIG. 3B.

[00106] In one embodiment, the web-based Program Neighborhood application includes a group of objects that manage various aspects of the application. In one embodiment, the application includes three primary object classes that "plug in" to a web server: a gateway object class; a credentials object class; and an applications object class. In some specific embodiments, the object classes are provided as Java beans. The three primary object classes facilitate: validation of user credentials into a server farm; generation of lists of published applications that a specified user may access; provision of detailed information about a specific published application; and conversion of published application information into an ICA-compatible format.

[00107] When provided as Java beans, the objects can be accessed in a number of different ways. For example, they may be compiled as COM objects and made available to the web server as ActiveX components. In another embodiment, the Java beans can be used in their native form, such as when the server uses Java Server Pages technology. In yet another embodiment,

- 28 -

the Java beans can be instantiated and used directly in a Java servlet. In still another embodiment, the server 30 can instantiate the Java beans as COM objects directly.

[00108] A credentials object class manages information necessary to authenticate a user into a target server farm. A credentials object passes stored user credentials to other Program Neighborhood objects. In some embodiments, the credentials object is an abstract class that cannot be instantiated and represents a user's credentials. Various class extensions may be provided to allow different authentication mechanisms to be used, including biometrics, smart cards, token-based authentication mechanisms such as challenge-response and time-based password generation, or others. For example, a "clear text credentials" extension may be provided that stores a user's name, domain, and password in plain text.

[00109] A gateway object class handles communications with a target server farm. In one embodiment, the gateway object class is provided as an abstract Java class that cannot be instantiated. A particular gateway object may retrieve application information by communicating with a server farm using a particular protocol, reading cached application information, a combination of these two methods, or other various methods.

[00110] As noted above, the gateway object class caches information to minimize communication with a target server farm. Extensions to the gateway object may be provided to communicate with the server farm over specific protocols, such as HTTP. In one embodiment, an extension class is provided that allows the gateway object to communicate with the server farm via WINDOWS NT named pipes. The gateway object may provide an application programming interface hook that allows other Program Neighborhood objects to query the object for application information.

[00111] An applications object class contains information about published applications and returns information about applications hosted by the server farm in order to create the Program Neighborhood web page. The applications object class creates objects representing applications by retrieving information relating to the applications, either from an object created by the gateway object or directly from the servers in the server farm. An applications object acts as a container for certain properties of the application, some settable and some not settable, such as: the name of the application (not settable); the percentage of the client's desktop that the client window should occupy (settable); the width of the client window, in pixels, for this application

- 29 -

(settable); the height off the client window, in pixels, for this application (settable); the number of colors to use when connecting to the application (settable); the severity of audio bandwidth restriction (settable); the level of encryption to use when connecting to the application (settable); the level of video to use when connecting to this application (settable); whether the application should be placed on a client's start menu (settable); whether the application should be placed on the client's desktop (settable); the identity of the Program Neighborhood folder to which the application belongs (settable); the description of the application (settable); the source of the graphics icon file for the application (settable); the type of window that should be used when connecting to the application (not settable); and whether to override default parameters for the object.

[00112] FIG. 6A is a screenshot of an exemplary Program Neighborhood window 120 that can be displayed on the screen 18, 22 of either client node 10, 20 after the Program Neighborhood application has executed. The window 120 includes graphical icons 122. Each icon 122 represents an application program that is hosted by one of the servers 30, 32, 34 and 36 on the network 40. Each represented application is available to the user of the client node for execution. The user can select and launch one of the applications using the mouse 18, 28 or keyboard 14, 24.

[00113] FIG. 6B is a screenshot of another exemplary Program Neighborhood window 124 that can be displayed on the screen 18, 22 of either client node 10, 20 after the Program Neighborhood application has executed. The window 124 includes graphical icons 126, 128. Each icon 126, 128 represents an application program that is hosted by one of the servers 30, 32, 34 and 36 on the network 40. Each application program represented by one of the icons 126 is available to the user of the client node 10, 20 for execution. The user can select and launch one of the applications using the mouse 18, 28 or keyboard 14, 24. For web-based program neighborhood environments, the screenshots of Figs. 6A and 6B are similar, except that icons 122, 166, 128 are displayed within a browser window.

[00114] Each application program represented by one of the icons 128 is unavailable to the user of the client node 10, 20, although such applications are present in the server farm. The unavailability of these application programs can be noted on the display screen (e.g., "X"s can be

- 30 -

drawn through the icons 128). An attempt to launch such an application program can trigger a message indicating that the user is not authorized to use the application.

[00115] FIG. 7 shows an exemplary process by which a user of either client node 10, 20 can be informed about the availability of applications hosted by application servers 30, 32, 34 and 36 on the network 40. In step 80, the client node 10, 20 requests log-on service from one of the servers, e.g., server 32. The server 32 requires valid user credentials to establish the connection. The server 32 receives the user credentials from the client node 10 (step 82), and authenticates the user for log-on (step 84). A desktop is displayed at the client node 10, 20 (step 85). The desktop can include a graphical icon representing the Program Neighborhood application program.

[00116] In step 86, the application server 32 establishes a connection with each other servers 30, 34 and 36 to exchange application-related information, as described above, corresponding to application programs hosted on those servers (step 88). In step 90, the application server 32 develops and maintains a database of the collected application-related information. Each other server 30, 34, 36 in the server farm can develop a database equivalent to the database of the server 32 and in similar manner as the server 32. In another embodiment, the database of the server 32 can be a centralized database that is accessible to each other application server 30, 34, 36 in the server farm. The collecting of application-related information can occur independently or be triggered by the request of the client node 10, 20 to log-on to the server farm 38.

[00117] In step 92, the client node 10, 20 can request execution of an application program from the desktop display. The master server node can process the request and, using a load-balancing evaluation and application availability as described above, determine the application server to provide the service to the client node 10, 20 (step 94). For example, the application server 32 can be selected to service the request with the client node 10, 20. In step 96, the client node 10, 20 establishes a communications link with the server 32. The server 32 and the client node 10, 20 can communicate according to the ICA protocol appropriate for that client node as described above.

[00118] Also in response to this request to run the application program, the master server node 30 or the server 32 can run the Program Neighborhood application (step 93) and push the results to the client node 10, 20, although the client node 10, 20 may not have requested the PN

- 31 -

application program. When executing the PN application program, the master server node 30 or server 32 filters the application-related information in the database using the user credentials (step 100). The result of filtering the database determines those application programs that are authorized for use by the user of the client node 10, 20. The authorized application programs are in the program neighborhood of the client node 10, 20. This program neighborhood of available application information is pushed to the client node 10, 20 (step 102) and displayed on the client screen 12, 22 in a Program Neighborhood graphical window (step 104).

[00119] In other embodiments, the Program Neighborhood window may include applications that are in the server farm but unavailable for use to the client node 10, 20. In a Windows-based implementation, the available (and unavailable) application programs can be represented by icons. The user of the client node 10, 20 can select and launch one of the application programs displayed in the Program Neighborhood window.

FILE-TYPE ASSOCIATION (FTA)

[00120] As discussed, in one embodiment the Program Neighborhood mechanism permits users to identify, select, and execute executable programs not necessarily installed on their client node. In another embodiment, additional FTA functionality permits users to automatically initiate the execution of executable programs associated with a data file, even though the data file and the executable program are hosted on different computing nodes.

[00121] Typically, FTA functionality permits users to transparently execute executable programs by selecting data files located on a computing node that differs from the node(s) where the executable programs are located. In one embodiment, a user can transparently invoke the execution of an executable program on a server node by selecting a data file located on their client node. In another embodiment, a user can transparently invoke the execution of an application program on their client node by selecting a data file located on a server node. In still another embodiment, a user can select a data file stored on a web server and transparently invoke the execution of an associated executable program on a server node. Typically, execution permits processing of the contents of the selected data file, the output of which is then provided to the user at the client node.

- 32 -

[00122] As described above, the client node is in communication with one or more server nodes. The client node executes an operating system that includes functionalities that permit the selection of graphical indicia representative of files, the display of bitmapped graphics, and the initialization and operation of telecommunications links. In one embodiment, these functionalities are integrated with the operating system. In another embodiment, these functionalities are provided by software from one or more third-party vendors that is separately installed and maintained on the client node.

[00123] The server node executes an operating system including functionality that permits the initialization and operation of telecommunications links and the execution of executable programs. In one embodiment, these functionalities are integrated with the operating system. In another embodiment, these functionalities are provided by software from one or more third-party vendors that is separately installed and maintained on the server node.

[00124] In the present discussion, it is to be understood that examples using filename extensions necessarily reflect the idiosyncrasies of embodiments utilizing the WINDOWS family of operating systems. Other embodiments implement methods and apparatus in accord with the present invention using special parameters stored in the data file itself, the data contained in the data file, the file system records associated with the data file, or a separate data file or database. For example, in embodiments utilizing the MacOS family of operating systems, the present invention utilizes file and application creator types and stores file-type association data in the Desktop file associated with each storage device. Embodiments using a UNIX-variant operating system utilize file extensions, embedded parameters, or other mechanisms as appropriate. Accordingly, the scope of the claims should not be read to be limited to embodiments relying on filename extensions or embodiments utilizing WINDOWS operating systems.

Client-Based FTA

[00125] Referring to FIG. 8A, in one embodiment the present invention enables transparent distributed program execution on a server node through the selection of graphical indicia representative of a data file located on the client node. The client node receives, from one of a plurality of server nodes, a mapping specifying an association between a type of data file and an executable program for execution on one of a plurality of server nodes (Step 106). The client node presents a graphical depiction of a data file stored on the client node (Step 114) and

- 33 -

receives a selection of the graphical depiction of the data file (Step 118). The client node identifies an executable program associated with the type of the selected data file using the received mapping (Step 122) and sends a request to execute the identified executable program (Step 126). In another embodiment, the client node initiates the execution of a local display application (Step 130) to receive application output data from the executing program (Step 134), which it displays to the end user (Step 138).

[00126] Still referring to FIG. 8A, when the client node receives the mapping (Step 106), the mapping may be received by itself, with several other mappings, or with other messages or data such as software updates. Table 3 illustrates a exemplary mapping provided in one embodiment of the invention:

TABLE 3

| File type: | Executable program: |
|----------------|---------------------|
| “.DOC”, “.RTF” | MSWORD.EXE |
| “.PDF” | ACROBAT.EXE |

[00127] In one embodiment, the mapping identifies an association between a particular executable program for use with a particular data file stored on the user's client node. For example, the mapping would specify the use of WORD to edit the file "RECORDS.DOC". In another embodiment, the mapping specifies an association between a particular executable program and a particular type of data file stored on the user's client node. For example, the mapping would specify the use of WORD to edit all files ending in the extension ".RTF". In yet another embodiment, the mapping specifies a particular type of executable file for use with a particular data file or a particular type of data file. For example, the mapping would specify the use of a "word processing application" when a data file containing word processing information is selected, as determined by examining the file or special parameters coded into or associated with the file.

[00128] In still another embodiment, the mapping specifies the relationship between an executable program and a data file in terms of a client node application that launches the executable program on a server node and displays the output from execution at the client node. For example, the mapping could specify that when a ".DOC" file is selected, the client node is to execute METAFRAME from Citrix Software of Ft. Lauderdale, FL, which in turn sends a request to one of a plurality of server nodes to execute WORD, receiving the output data from

- 34 -

execution for display to the user at the client node. In yet another embodiment, the entry for the associated executable program includes one or more parameters, flags, or arguments that control execution on the server node. For example, the entry specifying the local execution of METAFRAME and the remote execution of WORD can take the form

5 "C:\PROGRA~1\MF.EXE MSWORD.EXE" which indicates that the METAFRAME application is stored locally at C:\PROGRA~1\MF.EXE and the application to be executed on the server node is MSWORD.EXE.

[00129] In one embodiment, the executable program in the mapping is specified as a vendor name, e.g. "MICROSOFT WORD." Upon receipt, the client node identifies at least one of the

10 file path and the file name associated with the vendor name "MICROSOFT WORD" for later use when applying FTA functionality. In another embodiment, the executable program in the mapping is specified as a file name in accord with the naming convention utilized by the client node operating system. For example, on a client node employing a WINDOWS-family operating system, the associated executable program could be identified as "MSWORD.EXE." After

15 receiving this mapping, the client node either locally identifies and stores the path to the executable file, e.g., "C:\PROGRA~1\MSOFFICE\MSWORD.EXE", or it dynamically locates the executable file each time FTA is invoked for that particular executable program.

[00130] In one embodiment, the received mapping is incorporated into a file containing a system database that tracks associations between executable programs and data files or types of

20 data files. For example, the file could be the registry on a WINDOWS system or the desktop file on a MacOS system. In another embodiment, the mapping is stored in a file that is accessed by one or more executable programs in order to provide the aforementioned FTA functionality. For example, a Java-based executable could receive the mapping and store it in a text file on the client node for use by other Java-based executables. In yet another embodiment where no file

25 system functionality is present, for example, an embedded device, the mapping is received and stored in random-access memory at a specific address. In still another embodiment, the mappings are stored remotely on, for example, one of a plurality of servers and retrieved on an as-needed basis by executable programs on the client node. The mappings themselves can be stored in any data structure, including but not limited to an array, a linked list, a flat database, or

30 a relational database.

- 35 -

[00131] In one embodiment, these mappings are periodically updated using client-based, i.e., “pull” techniques. In another embodiment, these mappings are periodically updated using server-side, i.e., “push” techniques. In yet another embodiment, these mappings are updated using either push or pull techniques on an as needed basis, for example, when a system administrator
5 installs a new application program and provides support for it on a server node or across a server farm.

[00132] In one embodiment, the client node receives rules that control the location of execution of the executable program (Step 110). In one embodiment, the rules for determining whether a particular executable program executes on the client node, a server node, or a
10 particular server node may depend on one or more variables. Such variables include but are not limited to the execution load on the client system, the execution load on one of the plurality of server systems, the execution load on the collective plurality of server systems, the speed of the connection between the client system and the plurality of server systems, or any combination of these or other factors. If, for example, the mapping required that the selection of a “.DOC”-type
15 document initiate the execution of a local thin-client application that would receive and display output from WORD executing on a server node in a server farm, then a conditional rule could specify that if the execution load on the client system consumed more than 50% of available CPU time and the connection to the server farm either exceeded 1 Mbit/sec in theoretical capacity or 20 Kbit/sec in available capacity, then WORD should be executed remotely on the
20 server farm with its application output data provided to the thin-client application for display to the end user. In one embodiment, when the evaluation of a rule indicates that the associated executable program should execute on the client node, the client node provides the output of the executing program directly to the user, bypassing any thin-client application.

[00133] In actual operation, when the user selected a “.DOC” document, the associated
25 application, i.e., WORD, would be identified, any accompanying rules would be loaded and evaluated, and the associated application would either be executed locally on the client node or remotely on a server node in response to the evaluation of the rule. If the evaluation indicates the application program is to execute on a server node, a request to execute the program is transmitted to the server node (Step 126).

- 36 -

[00134] The receipt of mappings, conditional rules, or some combination of the two can occur through several mechanisms. Mappings and conditional rules may be distributed over a live telecommunications link, or off-line using one or more forms of portable storage media. They may be distributed in real-time, while the client node and at least one of the plurality of server nodes is in use, or off-line, either at a scheduled time, within a scheduled time window, or when either or both of the client node and at least one of the plurality of server nodes is not in active use by an end user.

[00135] With the mappings and, in some embodiments, conditional rules installed, the user interacts with the client node using its graphical user interface. The data files present on the local client node are presented graphically to the user (Step 114). In a typical embodiment, the client node displays one or more icons representative of data files stored on the client node. In another embodiment, the client node displays a list of file names associated with data files stored on the client node. In still another embodiment, indicia representative of files stored on the client node are intermingled with indicia representative of files stored on one or more server nodes. In this embodiment, client-based FTA is operative when indicia representative of a file stored on the client node is selected. In another embodiment, multiple forms of FTA (see below) are operative, with the appropriate form of FTA activated based on the location of the file associated with the selected indicia.

[00136] The user selects a particular data file by interacting with the graphical depiction of the data file (Step 118). The icon or text representative of the data file is typically selected using a keyboard, a mouse, a light pen, a pressure-sensitive tablet, a haptic interface, or other input device.

[00137] Once the data file is selected, functionality on the client node accesses the stored mappings, and in some embodiments, the stored conditional rules, and identifies those mappings and rules associated with the selected data file or the type of the selected data file (Step 122). If the mapping indicates that the data file is associated with an application to be executed on one of a plurality of servers, and the evaluation of any associated conditional rules so indicates, then the client node sends an execution request to one of a plurality of servers hosting the associated application program (Step 126).

- 37 -

[00138] Upon execution, the associated application program will typically generate one or more forms of output data, including but not limited to graphical output data or sound. In one embodiment, the client node initiates the execution of an application program (Step 130) that, through a connection with the server node, receives the application output data (Step 134) and displays it to the end user (Step 138). In one embodiment, this application program is METAFRAME.

[00139] In one embodiment, the client node also provides data associated with the selected data file to the server node for processing. In one embodiment, the client node utilizes a network connection with the server node to provide the data. In another embodiment, the client node provides a pointer to the contents of the data file to the server node. In turn, the server node generates sequential or random access requests for data utilizing the pointer, which it provides to the client node. The client node processes these requests and provides the appropriate data to the server node.

[00140] This client-server arrangement permits the usage of client nodes and server nodes based on different computing architectures. For example, the client node may be a POWERPC based MACINTOSH executing a local display application and the server node may be an ITANIUM based WINDOWS NT server. As long as the client node and the server node are equipped with functionalities that agree on the formats for the transmission, exchange, and display of application output data, binary-level interoperability between the nodes is unnecessary.

[00141] FIG. 8B illustrates the typical role of a server node in the client-based file-type association process. A mapping is provided specifying an association between a type of data file stored on a client system and an executable program for execution on one of a plurality of server nodes (Step 154). A request to execute the executable program is received (Step 162) and the executable program is executed on one of a plurality of server nodes (Step 166).

[00142] In one embodiment, a mapping is received at a server node from a system administrator or another user with administrative privileges (Step 150). The administrator configures a new mapping by editing a system database that specifies an association between a particular data file or type of data file and a particular executable program or type of executable program, as discussed above. In another embodiment, the administrator specifies an association between a type of data file, a program for establishing a connection with a server system, and

- 38 -

executable program for processing the data file, and parameters for use by the executable program. In one embodiment, the administrator specifies the mapping by manually editing one or more data files or by using a graphical system administration tool.

[00143] In one embodiment, the server node provides this system database in a file to one or more client nodes either periodically or on an as-needed basis using push or pull techniques. In another embodiment, these mappings are sent as standalone data entries to client systems that receive and incorporate them into their own local system databases. In still another embodiment, the specified mapping is also replicated or synchronized against the mappings specified on other server nodes. This last functionality facilitates the configuration and maintenance of server nodes in a server farm configuration.

[00144] The server node provides the mapping to one or more client nodes (Step 154) that utilize it as described above. In one embodiment, the server also provides conditional execution rules to one or more client nodes (Step 158) for storage or evaluation as described above. Conditional rules are specified and distributed using the techniques used for the specification and distribution of file-type associations.

[00145] In operation, the server nodes receive requests to execute an executable program from one or more client nodes (Step 162). In one embodiment, the availability of one or more executable programs on the server node(s) is determined in part by identification or authentication indicia. Each file on the server node is directly or indirectly associated with indicia that define a user, a group of users, or a privilege level. The client node provides one or more indicia indicative of the user's identity or the user's privilege level to a master server node or to one or more individual server nodes. The server node(s) verify these indicia against the indicia associated with each executable program available for execution. If the provided indicia match the associated indicia, then the executable program is made available for execution to the client node. Attempts to execute an executable program with authentication indicia exceeding those provided by the client node can result in one or more of an error message display on the client node, an alarm message display on a server node, and a request to initiate execution of the executable program on the client node.

[00146] The server node services these requests by executing the identified programs (Step 166). When servicing a request to execute a program, the server may itself service the request or

- 39 -

it may delegate or reassign the execution request to another server. This delegation facilitates license management and reissue by limiting the number of server licenses needed to serve a client pool. The server receiving the request may also assign the execution request to a computing cluster, enabling load and resource management among servers to better manage client demand.

[00147] In one embodiment, the server node also receives data associated with a particular data file identified on the client node (Step 170). Typically, this data is processed by the executable program as it executes on a server node or cluster as described above. In another embodiment, the server receives pointers or other references to a data file on the client node that the server utilizes to generate requests for data, e.g., sequential or random access requests, for transmission to the appropriate client node. In one embodiment, if a first server node delegates a program execution request to a second server node, then the first server node, if it receives data or data file references, may similarly delegate, reroute or retransmit the data or data file references to the second server node or another server node, as appropriate.

[00148] Typically the server node provides the application output data from the execution of the executable program to the client node for display to the end user, as described above (Step 174). Application output data may be provided in one or more formats supported either directly by the client node, or by an application program executing on the client node, e.g., a thin-client display application such as METAFRAME. Application output data may be a bitstream representative of graphical data, compressed or uncompressed; a markup language representative of vector graphics for display; text; or a series of system-dependent or API calls that when executed display the application output data on the client node for the user.

Server-Based FTA

[00149] Referring to FIG. 9A, in another embodiment the present invention enables transparent distributed program execution on a client node through the selection of graphical indicia representative of a data file located on a server node. The client node presents a graphical depiction of a data file stored on one of a plurality of server nodes (Step 200). The client node receives a selection of the graphical depiction of the data file (Step 204) and transmits the selection to one of the plurality of server nodes (Step 208). The client node receives a request

- 40 -

from one of the plurality of server nodes to execute an executable program associated with the selected data file (Step 212) and executes the associated executable program (Step 216).

[00150] Still referring to FIG. 9A, the client node presents a user with a graphical depiction of at least one data file stored on at least one server node (Step 200). The graphical depiction of the data files includes but is not limited to the display of graphical indicia such as icons or the display of file names. In one embodiment, indicia representative of files stored on one or more server nodes are intermingled with indicia representative of files stored on the client node. In this embodiment, server-based FTA is operative when indicia representative of a file stored on a server node is selected. In another embodiment, multiple forms of FTA (see above, below) are operative, with the appropriate form of FTA activated based on the location of the file associated with the selected graphical indicia.

[00151] In one embodiment, functionality on the client node queries a master server node which aggregates information concerning the data files available on connected server nodes. The master server node presents an aggregated list of files available on the master server and connected servers and presents it to the client node using either push or pull mechanisms before presentation to the user. In another embodiment, functionality on the client node queries each individual server node concerning the availability of data files for server-side FTA before presenting available files to the user. In yet another embodiment, functionality on the client node (e.g., Network File System (NFS) or Andrew File System (AFS)) makes the files on one or more server nodes appear as part of the client node's file system. These files are subsequently presented to the user for use with server-based FTA.

[00152] In one embodiment, the availability of one or more files on the server node(s) is determined in part by identification or authentication indicia. Each file on the server node is directly or indirectly associated with indicia that define a user, a group of users, or a privilege level. The client node provides one or more indicia indicative of the user's identity or the user's privilege level to a master server node or to one or more individual server nodes. The server node(s) verify these indicia against the indicia of each file available for server-based FTA. If the provided indicia match the associated indicia, then the file is made available to the client node and graphically presented to the user for potential server-based FTA activities.

- 41 -

[00153] As described above, the user may select a data file by interacting with the displayed graphical indicia, typically using a graphical user interface and one or more input devices (Step 204). Once the user selects a data file, the selection is transmitted to one of the plurality of server systems (Step 208). In one embodiment, the selection is transmitted to a master server node.

5 The master server node identifies the server node hosting the data file and provides the selection to the appropriate server node for further processing. In another embodiment, the selection is transmitted to one or more server nodes, either server nodes specifically associated with the selected data file or server nodes identified as generally available for server-based FTA. The latter behavior facilitates load-balanced file serving among servers through file mirroring and
10 other advanced forms of data redundancy.

[00154] After providing the selection of the data file, a request to execute an executable program associated with the selected data file is received (Step 212). As discussed above, the request may come from a master server node, a server node hosting the selected data file, or a server node not hosting the selected data file. In response, the client node executes the
15 executable program associated with the data file (Step 216).

[00155] In one embodiment, the client node receives data associated with the selected data file for processing by the executing program (Step 220). In another embodiment, the client node receives a pointer to the selected data file. The client node utilizes the provided pointer to generate one or more requests for data, e.g. sequential or random, which it provides to one or
20 more server nodes for servicing by a master server node, a server node hosting the selected data file, or a server node not hosting the selected data as discussed above. In one embodiment, these requests are transmitted utilizing a protocol, such as hypertext transfer protocol (HTTP), secure HTTP (HTTPS) (i.e., using SSL), messaging application programming interface (MAPI) protocol, real-time streaming protocol (RTSP), real-time streaming protocol used for user
25 datagram protocol scheme (RTSPU), the Progressive Networks Multimedia (PNM) protocol developed by RealNetworks, Inc. of Seattle, WA, or manufacturing message specification (MMS) protocol.

[00156] Having received data associated with the selected data file, the client node typically processes the received data using the executing program (Step 224). The client node typically
30 displays the result of the processing to the end user. In one embodiment, the processing and

- 42 -

display cycle is iterative, with further control input from the user leading to further processing and display, leading in turn to further control input from the user, etc.

[00157] FIG. 9B illustrates the typical role of a server node in server-side FTA. A mapping is received that specifies an association between a type of data file and an executable program for execution on a client system (Step 240). A data file is stored on one of a plurality of server systems (Step 244). A selection of the stored data file is received by the server node (Step 248) and an executable program associated with the type of the selected data file is identified using the received mapping (Step 252). A request is sent to the client system to execute the identified executable program (Step 256).

[00158] The mapping received by the server node (Step 240) specifies an association between a data file and an executable program and, in this respect, is similar to the mapping utilized in client-based FTA. The mapping is different from the mapping used in client-based FTA in that it specifies an executable program for execution on the client node. For embodiments in which the application executes at the client node, the associated program is typically a standalone application and not a thin-client application that invokes the application of a second executable program.

[00159] However, in some embodiments, the client node receives the application execution request from the server node and, instead of executing the identified executable program (and thus executing the application), redirects the execution request to another server node. Thus, the server node to which the client node redirects the execution request is the node which executes the identified executable program. In the embodiments, the client-based FTA "overrides" the server-based FTA because the server node identified by the client-side mapping executes the application, not the node identified by the server-side mapping.

[00160] In one embodiment, the mapping is received as a stand-alone entry. In another embodiment, mappings are aggregated and presented in a group, or aggregated and transmitted with other data such as software updates. Mappings are provided over a live telecommunications link or stored in binary or text format on one or more removable storage media.

[00161] As discussed above, the mapping specifies a linkage between data files or executable programs at various levels of granularity. Particular mappings link particular data files or

- 43 -

particular types of data files and particular executable programs or types of executable programs. Executable programs are identified by vendor name or by filename.

5 [00162] In one embodiment, the received mapping is incorporated in a file containing a system database that tracks associations between executable programs and data files or types of data files. In another embodiment, these mappings are updated on either a periodic or on an as-needed basis using push techniques, pull techniques, or a combination of the two.

10 [00163] In still another embodiment, the server system receives one or more rules for determining whether an identified associated executable program is to be executed on the client system or one of the plurality of server systems. These rules are functions of one or more variables including but not limited to the execution load on the client system, the execution load on one of the plurality of server systems, the execution load on the collective plurality of server systems, the speed of the connection between the client system and the plurality of server systems, or any combination of these or other factors.

15 [00164] In operation, one or more data files are stored on one or more server nodes (Step 244). Data files are provided to server nodes over a live telecommunications link or off-line using one or more forms of portable storage media. Data files are provided at a scheduled time, within a scheduled time window, or when the client node in question is idle or has a processor execution load below a certain predetermined level.

20 [00165] The server node receives a selection of graphical indicia representative of the stored data file (Step 248). When the selection is received, the server identifies any associated executable program or type of executable program identified in the received mapping (Step 252). A request to execute the identified executable program is transmitted to the client node for execution (Step 256).

25 [00166] In one embodiment, any associated rules are loaded and evaluated to determine whether the executable program is to execute on the client node, one of the plurality of server nodes, or several server nodes as in a distributed processing environment. If the evaluation of the rule indicates that execution is to occur on the client node, then a request to execute the associated executable program is transmitted to the client node (Step 256). In one embodiment, when the evaluation of a rule indicates that the associated executable program should execute on

- 44 -

the client node, the client node provides the output of the executing program directly to the user, bypassing any thin-client application in operation on the client node.

[00167] In one embodiment, a server node provides data associated with the selected data file for processing (Step 260). Typically this information is processed by the executable program operating on the client node and the result is provided to the user, where further input cues further processing and display. In another embodiment, the server node provides a pointer or other reference to the selected data file. The client node generates requests for data, e.g., random or sequential, using the provided reference. The requests are transmitted to the server node, which services them by providing data in response to the request.

10 [00168] In yet another embodiment, services provided by a server node, e.g., servicing data requests or providing data associated with a selected file, are provided by a plurality of server nodes acting in concert, e.g., a server farm. In one embodiment a master server node receives requests for processing or data and services them directly. In another embodiment, the master server node delegates client requests to other server nodes. This technique facilitates balanced processing among server nodes. In yet another embodiment, requests are subdivided and distributed among several server nodes that handle processing or data access in parallel.

Web-Based FTA

[00169] Referring to FIG. 10A, in one embodiment the present invention enables transparent program execution on a server node through the selection of graphical indicia presented on a client node that are representative of data files located on a web server. The client node presents a graphical depiction of a data file stored on a web server (Step 280). The client node receives a selection of a graphical depiction of a data file (Step 284) and provides it to the web server (Step 288). The client node receives a request to execute a first executable program, the request including information concerning a second executable program associated with the data file for execution on one of the client system and one of a plurality of server systems (Step 292). The client node executes the first executable program (Step 296) and provides the request to execute the second executable program (Step 300). The client node receives application output data from the execution of the second executable program (Step 304) and displays it to the end user (Step 308).

- 45 -

[00170] Still referring to FIG. 10A, the client node presents a user with a graphical depiction of one or more data files stored on one or more web servers (Step 280). The graphical depiction of the data files includes but is not limited to the display of graphical indicia such as icons or the display of file names as hyperlinks in an SGML document. In one embodiment, indicia representative of files stored on web servers are intermingled with indicia representative of files stored on the client node or files stored on one or more server nodes. In this embodiment, web-based FTA operates when an icon representative of a file stored on a web server is selected. In another embodiment, multiple forms of FTA (see above) are operative, with the appropriate form of FTA activated based on the location of the file associated with the selected graphical indicia.

[00171] The client node receives the user's selection of a particular data file using a keyboard, a mouse, a light pen, a pressure-sensitive tablet, a haptic interface, or other input device (Step 284). The client node provides the selection to the web server (Step 288), typically using a telecommunications protocol such as HTTP.

[00172] In response to providing the selected data file, the client node receives an execution request (Step 292). The request includes several parameters including but not limited to a reference to a first executable program and a reference to a second executable program associated with the data file. In one embodiment, the executable programs are identified by vendor names or file names which are resolved by the client system into at least one of a file name, a path name, or a reference to the desired executable program. The request typically includes a parameter that identifies the selected data file on the web server, such as a Uniform Resource Locator (URL).

[00173] In one embodiment, an exemplary request is "MetaFrame MSWORD.EXE http://12.34.127.128/doc/VendorList.RTF." The first parameter is "MetaFrame," representative of the first executable program for execution on the client node, which the client node identifies as the executable program stored at "C:\PROGRA~1\CITRIX\MF.EXE". The second parameter is "MSWORD.EXE," which refers to the second executable program for resolution and execution on a server node. The third parameter is "http://12.34.127.128/doc/VendorList.RTF" which identifies the file on the web server that the user had selected. In operation, the client node executes MetaFrame and forwards the request to execute WORD to a server node.

- 46 -

[00174] In further embodiments, the server node transmits a request which includes control information, such as a security token, to the client node. For example, the first executable program may include control information in the request to execute the second executable program on a server node. Thus, the control information may invoke a third program on the server node or web server upon receipt of a request to execute the second executable program (e.g., WORD). This third program could modify the manner of delivery of the application (e.g., require the user to provide appropriate authentication credentials, such as a password, before delivery, require the server node to wait for a key from another node before delivery of the application). In a further embodiment, the control information may provide a server node with an indicia to enable the server node to access a second executable program previously inaccessible by the server node.

[00175] In one embodiment, the selected file on the web server is received by the client node, which caches the file for editing. In this embodiment, the client node provides data associated with the file to the server node for processing by the executing program. In another embodiment, the client node provides a reference to the cached file to the server node. The server node uses the provided reference to generate requests for data which the client node services. In still another embodiment, the parameter identifying the selected data file is provided directly to the server node, which contacts the web site directly for data associated with the selected file or delegates the request to another server in a server farm.

[00176] The client node executes the first executable program (Step 296) which, in one embodiment, is a thin-client application like METAFRAME for receiving and displaying application output data to an end user. In one embodiment, the client node receives the first executable program from a server node or a web server using push techniques, pull techniques, or a combination thereof on either an as-needed or a scheduled basis.

[00177] The client node provides the request to execute the second executable program to a master server node, a server node, or a plurality of server nodes (Step 300). In one embodiment, a plurality of server nodes coordinate and share the execution task. This functionality permits the implementation of load balancing in a server farm or server array.

[00178] In one embodiment, the client node receives one or more rules for determining whether the second executable program is to be executed on the client system or one of the

- 47 -

plurality of server systems. These rules are functions of one or more variables including but not limited to the execution load on the client system, the execution load on one of the plurality of server systems, the execution load on the collective plurality of server systems, the speed of the connection between the client system and the plurality of server systems, or any combination of these or other factors. In operation, any associated rules are loaded and evaluated to determine whether the second executable program is to execute on the client node, one of the plurality of server nodes, or several server nodes as in a distributed processing environment. In one embodiment, if the evaluation of the rule indicates that execution is to occur on the client node, the client node provides the output of the executing program directly to the user, bypassing any thin-client application in operation on the client node.

[00179] Upon execution, the associated application program will typically generate one or more forms of output data, including but not limited to graphical output data or sound. In one embodiment, the client node, through a connection with the server node, receives the application output data (Step 304) and displays it to the end user (Step 308). In one embodiment, application output data is provided in one or more formats supported either directly by the client node, or by the first executable program or another executable program executing on the client node. In another embodiment, application output data is one or more of a bitstream representative of graphical data, compressed or uncompressed; a markup language representative of vector graphics for display; text; or a series of system-dependent or API calls that when executed display the application output data on the client node for the user.

[00180] FIG. 10B illustrates the role of the web server in web-based FTA. The web server provides a graphical depiction of a data file (Step 320). The web server subsequently transmits an execution request to the client node (Step 324).

[00181] In one embodiment, the web server provides graphical indicia related to data files that it hosts that are available for selection by the end user (Step 320). In another embodiment, the web server determines which data files are available to the end user by comparing authentication or identity indicia against indicia associated with each hosted file. In still another embodiment, the web server prompts other connected server nodes or web servers to provide graphical indicia related to the files that they host which are accessible to the user, either aggregating the indicia and presenting them itself or leaving each individual server responsible for presentation.

- 48 -

[00182] The web server typically receives a selection of indicia associated with a presented file. In embodiments where the selected file is associated with a connected server, the receiving web server may either delegate the received selection to the appropriate server or retrieve the data associated with the selected file from the appropriate server, handling executable and data requests itself. In one embodiment, the choice between these options is made in accord with an evaluation designed to balance loads among servers.

[00183] In one embodiment, the web server receives individual mappings for incorporation into a data file for later use. In still another embodiment, the web server receives mappings in groups or aggregated with other data such as software updates. Mappings may be transmitted over live telecommunications links or stored in data files on one or more removable storage media. In one embodiment, these mappings are updated on either a periodic or on an as-needed basis using push techniques, pull techniques, or a combination of the two.

[00184] Utilizing the received mappings, the web server determines whether the selected data file is associated with an executable program for execution using web-based FTA mechanisms. The mappings typically specify an association between data files or executable programs at various levels of granularity. In one embodiment, mappings link particular data files or types of data files and particular executable programs or types of executable programs, identified either by vendor name or filename. In another embodiment, the mappings file also includes parameters related to user identity, including but not limited to one or more of a user name, a public key, and an IP address. This functionality permits the web server to maintain multiple sets of file-type associations that vary among individual users.

[00185] Using the received mapping information, the web server constructs and transmits an execution request to the client node (Step 324), which the client node subsequently executes as described above. In one embodiment, the web server also provides the first executable program to the client node for execution. In another embodiment, the web server also provides a reference to the selected data file that enables other computing nodes to formulate requests for data associated with the selected data file. In still another embodiment, the web server also services data requests from other computing nodes formulated using the provided reference. In still another embodiment, the server node provides the data associated with the selected data file to the client node or one of a plurality of server nodes to facilitate subsequent processing.

- 49 -

[00186] In still another embodiment, the web server also provides one or more rules to the client node that permits the client node to decide whether the second executable program should be executed on the client node, one of the plurality of server nodes, or a plurality of server nodes based on factors including but not limited to the load on the client system, the load on a particular server system, the load on the plurality of server systems, or the available or theoretical bandwidth of the connection between one or more of the client node, the web server, a server node, or the plurality of server nodes.

[00187] The present invention may be provided as one or more computer-readable programs embodied on or in one or more articles of manufacture. The article of manufacture may be a floppy disk, a hard disk, a CD ROM, a flash memory card, a PROM, a RAM, a ROM, or a magnetic tape. In general, the computer-readable programs may be implemented in any programming language. Some examples of languages that can be used include C, C++, or JAVA. The software programs may be stored on or in one or more articles of manufacture as object code.

[00188] Many alterations and modifications may be made by those having ordinary skill in the art without departing from the spirit and scope of the invention. Therefore, it must be expressly understood that the illustrated embodiments have been shown only for the purposes of example and should not be taken as limiting the invention, which is defined by the following claims. These claims are to be read as including what they set forth literally and also those equivalent elements which are insubstantially different, even though not identical in other respects to what is shown and described in the above illustrations.

- 50 -

CLAIMS

What is claimed is:

- 1 1. In a network including a client system and a plurality of server systems, a method for
2 enabling distributed program execution, the method comprising the steps:
 - 3 (a) receiving, from one of a plurality of server systems, a mapping specifying an
4 association between a type of data file and an executable program for execution on
5 one of a plurality of server systems;
 - 6 (b) presenting a graphical depiction of a data file stored on a client system;
 - 7 (c) receiving a selection of the graphical depiction of the data file;
 - 8 (d) identifying an executable program associated with the type of the selected data file
9 using the received mapping; and
 - 10 (e) sending a request to execute the identified executable program.
- 1 2. The method of claim 1 wherein step (a) comprises receiving a mapping specifying an
2 association between a type of data file, a program for establishing a connection with a server
3 system, an executable program for processing the data file, and parameters for use by the
4 executable program.
- 1 3. The method of claim 1 further comprising the step:
2 modifying a file on the client system to include the received mapping.
- 1 4. The method of claim 1 further comprising the step:
2 executing a second executable program for displaying application output data.
- 1 5. The method of claim 1 further comprising the step:
2 receiving application output data from the executing program.
- 1 6. The method of claim 1 further comprising the step:
2 updating the received mapping on at least one of a periodic basis and an as-needed basis.
- 1 7. The method of claim 1 further comprising the step:
2 receiving a rule determining whether an identified executable program is to be executed
3 on the client system or one of the plurality of server systems.

- 51 -

- 1 8. The method of claim 1 further comprising the step:
2 receiving a rule for determining whether an identified executable program is to be
3 executed on the client system or one of the plurality of server systems that specifies
4 the determination between the client system and one of the plurality of server systems
5 as a function of at least one of the load on the client system, the load on one of the
6 plurality of server systems, the load on the plurality of server systems, and the speed
7 of the connection between the client system and the plurality of server systems.
- 1 9. The method of claim 1 wherein step (e) comprises the steps:
2 (e-a) receiving a rule determining whether an identified executable program is to be
3 executed on the client system or one of the plurality of server systems;
4 (e-b) evaluating the rule to determine whether the identified executable program is to be
5 executed on the client system or one of the plurality of server systems; and
6 (e-c) sending a request to execute the identified executable program on one of the
7 plurality of server systems based on the evaluation in step (e-b).
- 1 10. The method of claim 1 wherein step (e) comprises the steps:
2 (e-a) receiving a rule for determining whether an identified executable program is to be
3 executed on the client system or one of the plurality of server systems that specifies
4 the determination between the client system and one of the plurality of server systems
5 as a function of at least one of the load on the client system, the load on one of the
6 plurality of server systems, the load on the plurality of server systems, and the speed
7 of the connection between the client system and the plurality of server systems;
8 (e-b) evaluating the rule to determine whether the identified executable program is to be
9 executed on the client system or one of the plurality of server systems; and
10 (e-c) sending a request to execute the identified executable program on one of the
11 plurality of server systems based on the evaluation in step (e-a).
- 1 11. The method of claim 1 further comprising the step:
2 (f) transmitting data comprising the selected data file, at least some of the data for
3 processing by the executing program.
- 1 12. The method of claim 11 wherein step (f) comprises the steps:

- 52 -

- 2 (f-a) providing a pointer to the selected data file; and
3 (f-b) responding to requests for data associated with the selected data file utilizing the
4 provided pointer.
- 1 13. In a network including a client system and a plurality of server systems, a method for
2 enabling distributed program execution, the method comprising the steps:
3 (a) providing a mapping specifying an association between a type of data file stored on a
4 client system and an executable program for execution on one of a plurality of server
5 systems;
6 (b) receiving a request to execute the executable program; and
7 (c) executing the executable program on one of a plurality of server systems.
- 1 14. The method of claim 13 wherein step (a) comprises providing a mapping specifying an
2 association between a type of data file, a program for establishing a connection with a server
3 system, an executable program for processing the data file, and parameters for use by the
4 executable program.
- 1 15. The method of claim 13 further comprising the step:
2 providing application output data from the server system executing the executable
3 program to a client system.
- 1 16. The method of claim 13 further comprising the step:
2 providing a second application for displaying application output data from the server
3 system executing the executable program.
- 1 17. The method of claim 13 further comprising the step:
2 updating the provided mapping on at least one of a periodic basis and an as-needed basis.
- 1 18. The method of claim 13 further comprising the step:
2 (d) receiving data associated with a data file for processing by the executing program.
- 1 19. The method of claim 18 wherein step (d) comprises the steps:
2 (d-a) receiving a pointer to the data file; and
3 (d-b) sending requests for data associated with the data file utilizing the received pointer.

- 53 -

- 1 20. The method of claim 18 further comprising the step:
2 processing the received data using the executable program.
- 1 21. The method of claim 20 further comprising the step:
2 providing the processed data.
- 1 22. The method of claim 13 further comprising the step:
2 providing a rule to determine whether an executable program is to be executed on a client
3 system or one of a plurality of server systems.
- 1 23. The method of claim 13 further comprising the step:
2 providing a rule for determining whether an executable program is to be executed on the
3 client system or one of the plurality of server systems that specifies the determination
4 between the client system and one of the plurality of server systems as a function of at
5 least one of the load on the client system, the load on one of the plurality of server
6 systems, the load on the plurality of server systems, and the speed of the connection
7 between the client system and the plurality of server systems.
- 1 24. In a network including a client system and a plurality of server systems, a method for
2 enabling distributed program execution, the method comprising the steps:
3 (a) receiving, by the client system, a mapping specifying an association between a type of
4 data file stored on the client system and an executable program for execution on one
5 of a plurality of server systems;
6 (b) presenting, by the client system, a graphical depiction of a data file stored on the client
7 system;
8 (c) receiving, by the client system, a selection of the graphical depiction of the data file;
9 (d) identifying, by the client system, an executable program associated with the type of
10 the selected data file using the received mapping; and
11 (e) sending a request to execute the identified executable program to one of a plurality of
12 server systems.
- 1 25. The method of claim 24 wherein step (a) comprises receiving a mapping specifying an
2 association between a type of data file, a program for establishing a connection with a server

- 54 -

3 system, an executable program for processing the data file, and parameters for use by the
4 executable program.

1 26. The method of claim 24 further comprising the step:
2 modifying, by the client system, a file to include the received mapping.

1 27. The method of claim 24 further comprising the step:
2 providing, to the client system, a second executable program for displaying application
3 output.

1 28. The method of claim 24 further comprising the step:
2 updating, by one of the plurality of server systems, the received mapping on at least one
3 of a periodic basis and an as-needed basis.

1 29. The method of claim 24 further comprising the step
2 executing, by one of a plurality of server systems, the executable program.

1 30. The method of claim 29 further comprising the step:
2 providing, by the server system executing the executable program, application output
3 data.

1 31. The method of claim 24 further comprising the step:
2 (f) transmitting, by the client system, data comprising the selected data file, at least some
3 of the data for processing by the executing program.

1 32. The method of claim 31 wherein step (f) comprises the steps:
2 (f-a) providing, by the client system, a pointer to the selected data file; and
3 (f-b) responding, by the client system, to requests for data associated with the selected
4 data file utilizing the provided pointer.

1 33. The method of claim 31 further comprising the step:
2 processing, by one of the plurality of server systems, the provided data using the
3 executable program.

1 34. The method of claim 33 further comprising the step:
2 providing, by one of the plurality of server systems, the processed data.

- 55 -

- 1 35. The method of claim 24 further comprising the step:
2 receiving, at the client system, a rule determining whether an identified executable
3 program is to be executed on the client system or one of the plurality of server
4 systems.
- 1 36. The method of claim 24 further comprising the step:
2 receiving, at the client system, a rule for determining whether an identified executable
3 program is to be executed on the client system or one of the plurality of server
4 systems that specifies the determination between the client system and one of the
5 plurality of server systems as a function of at least one of the load on the client
6 system, the load on one of the plurality of server systems, the load on the plurality of
7 server systems, and the speed of the connection between the client system and the
8 plurality of server systems.
- 1 37. The method of claim 24 wherein step (e) comprises the steps:
2 (e-a) receiving a rule determining whether an identified executable program is to be
3 executed on the client system or one of the plurality of server systems;
4 (e-b) evaluating a rule to determine whether the identified executable program is to be
5 executed on the client system or one of the plurality of server systems; and
6 (e-c) sending a request to execute the identified executable program on one of the
7 plurality of server systems based on the evaluation in step (e-a).
- 1 38. The method of claim 24 wherein step (e) comprises the steps:
2 (e-a) receiving, at the client system, a rule for determining whether an identified
3 executable program is to be executed on the client system or one of the plurality of
4 server systems that specifies the determination between the client system and one of
5 the plurality of server systems as a function of at least one of the load on the client
6 system, the load on one of the plurality of server systems, the load on the plurality of
7 server systems, and the speed of the connection between the client system and the
8 plurality of server systems;
9 (e-b) evaluating the received rule to determine whether the identified executable program
10 is to be executed on the client system or one of the plurality of server systems; and

- 56 -

11 (e-c) sending a request to execute the identified executable program on one of the
12 plurality of server systems based on the evaluation in step (e-a).

1 39. In a network including a client system and a plurality of server systems, a method for
2 enabling distributed program execution, the method comprising the steps:
3 (a) presenting a graphical depiction of a data file stored on one of a plurality of server
4 systems;
5 (b) receiving a selection of the graphical depiction of the data file;
6 (c) transmitting, to one of the plurality of server systems, the selection of the graphical
7 depiction of the data file;
8 (d) receiving a request, from one of the plurality of server systems, to execute an
9 executable program associated with the selected data file; and
10 (e) executing the executable program associated with the selected data file.

1 40. The method of claim 39 further comprising the step:
2 (f) receiving data associated with the selected data file for processing by the executing
3 program.

1 41. The method of claim 40 wherein step (f) comprises the step:
2 (f-a) receiving a pointer to the selected data file.

1 42. The method of claim 41 wherein step (f) further comprises the step:
2 (f-b) retrieving data, from one of the plurality of server systems, associated with the
3 selected data file utilizing the received pointer.

1 43. The method of claim 40 wherein step (f) comprises the step:
2 (f-a) receiving data associated with the selected data file for processing using at least one
3 of hypertext transfer protocol (HTTP) and messaging application programming
4 interface (MAPI) protocol.

1 44. The method of claim 40 further comprising the step:
2 processing the received data using the executing program.

1 45. The method of claim 44 further comprising the step:
2 providing the processed data.

- 57 -

1 46. In a network including a client system and a plurality of server systems, a method for
2 enabling distributed program execution, the method comprising the steps:

- 3 (a) receiving a mapping specifying an association between a type of data file and an
4 executable program for execution on a client system;
5 (b) storing a data file on one of a plurality of server systems;
6 (c) receiving a selection of the stored data file;
7 (d) identifying an executable program associated with the type of the selected data file
8 using the received mapping; and
9 (e) sending a request to a client system to execute the identified executable program.

1 47. The method of claim 46 further comprising the step:
2 modifying a data file on one of the plurality of server systems to include the received
3 mapping.

1 48. The method of claim 46 further comprising the step:
2 updating the received mapping on at least one of a periodic basis and an as-needed basis.

1 49. The method of claim 46 further comprising the step:
2 receiving a rule for determining whether an identified executable program is to be
3 executed on the client system or one of the plurality of server systems.

1 50. The method of claim 46 further comprising the step:
2 receiving a rule for determining whether an identified executable program is to be
3 executed on the client system or one of the plurality of server systems that specifies
4 the determination between the client system and one of the plurality of server systems
5 as a function of at least one of the load on the client system, the load on one of the
6 plurality of server systems, the load on the plurality of server systems, and the speed
7 of the connection between the client system and the plurality of server systems.

1 51. The method of claim 46 wherein step (e) comprises the step:
2 (e-a) evaluating a rule to determine whether the identified executable program is to be
3 executed on a client system or one of the plurality of server systems.

1 52. The method of claim 51 wherein step (e) further comprises the step:

- 58 -

(e-b) sending a request to execute the identified executable program on the client system based on the evaluation in step (e-a).

53. The method of claim 46 further comprising the step:

(f) providing data associated with the selected data file for processing by the executing program.

54. The method of claim 53 wherein step (f) comprises the step:

(f-a) providing a pointer to the selected data file.

55. The method of claim 54 further comprising the step:

(f-b) responding to requests for data associated with the selected data file utilizing the provided pointer.

56. In a network including a client system and a plurality of server systems, a method for enabling distributed program execution, the method comprising the steps:

(a) receiving, by one of a plurality of server systems, a mapping specifying an association between a type of data file and an executable program for execution on a client system;

(b) storing, on one of a plurality of server systems, a data file ;

(c) receiving, at a client system, a selection of the stored data file;

(d) transmitting, to one of a plurality of server systems, the selection of the data file;

(e) identifying, at one of a plurality of server systems, an executable program associated with the type of the selected data file using the received mapping;

(f) sending, to the client system, a request to execute the identified executable program; and

(g) executing, at the client system, the executable program associated with the selected data file.

57. The method of claim 56 further comprising the step:

(h) retrieving, by the client system, data associated with the selected data file.

58. The method of claim 57 wherein step (h) comprises the step:

- 59 -

2 (h-a) retrieving data associated with the selected data file for processing using at least one
3 of hypertext transfer protocol (HTTP) and messaging application programming
4 interface (MAPI) protocol.

1 59. The method of claim 56 further comprising the step:
2 providing, to the client system, a pointer to the selected data file.

1 60. The method of claim 59 further comprising the step:
2 retrieving, by the client system, data associated with the selected data file using the
3 provided pointer.

1 61. The method of claim 56 further comprising the step:
2 receiving a rule for determining whether an identified executable program is to be
3 executed on the client system or one of the plurality of server systems.

1 62. The method of claim 56 further comprising the step:
2 receiving a rule for determining whether an identified executable program is to be
3 executed on the client system or one of the plurality of server systems that specifies
4 the determination between the client system and one of the plurality of server systems
5 as a function of at least one of the load on the client system, the load on one of the
6 plurality of server systems, the load on the plurality of server systems, and the speed
7 of the connection between the client system and the plurality of server systems.

1 63. The method of claim 56 wherein step (f) comprises the step:
2 (f-a) evaluating, by one of a plurality of server systems, a rule to determine whether the
3 identified executable program is to be executed on a client system or one of the
4 plurality of server systems.

1 64. The method of claim 63 wherein step (f) further comprises the step:
2 (f-b) sending a request to execute the identified executable program on the client system
3 based on the evaluation in step (f-a).

1 65. The method of claim 56 further comprising the step:
2 processing the received data using the executing program.

- 60 -

- 1 66. The method of claim 65 further comprising the step:
2 providing the processed data.
- 1 67. In a network including a client system and a plurality of server systems, a method for
2 enabling distributed program execution, the method comprising the steps:
3 (a) presenting a graphical depiction of a data file stored on a web server;
4 (b) receiving a request to execute a first executable program, said request including
5 information concerning a second executable program associated with the data file, the
6 second program executing on one of the client system and one of a plurality of server
7 systems;
8 (c) receiving application output data from the second executable program; and
9 (d) formatting, by the first executable program, received application output data for
10 presentation.
- 1 68. The method of claim 67 further comprising the step:
2 receiving the first executable program.
- 1 69. The method of claim 67 further comprising the step:
2 receiving a rule for determining whether an identified executable program is to be
3 executed on the client system or one of the plurality of server systems.
- 1 70. The method of claim 67 further comprising the step:
2 receiving a rule for determining whether an identified executable program is to be
3 executed on the client system or one of the plurality of server systems that specifies
4 the determination between the client system and one of the plurality of server systems
5 as a function of at least one of the load on the client system, the load on one of the
6 plurality of server systems, the load on the plurality of server systems, and the speed
7 of the connection between the client system and the plurality of server systems.
- 1 71. The method of claim 67 further comprising the step:
2 (e) sending a request to execute the second executable program associated with the
3 selected data file.
- 1 72. The method of claim 71 wherein step (e) comprises the steps:

- 61 -

- 2 (e-a) evaluating a rule to determine whether the second executable program is to be
3 executed on the client system or one of the plurality of server systems; and
4 (e-b) sending a request to execute the second executable program on one of the plurality
5 of server systems based on the evaluation in step (e-a).
- 1 73. The method of claim 67 further comprising the step:
2 receiving application output data from the second executable program.
- 1 74. The method of claim 73 further comprising the step:
2 displaying the application output data using the first executable program.
- 1 75. The method of claim 67 further comprising the step:
2 receiving data associated with the data file.
- 1 76. The method of claim 75 further comprising the step:
2 providing the received data to the computer executing the second executable program.
- 1 77. The method of claim 67 further comprising the step:
2 receiving a pointer to the data file.
- 1 78. The method of claim 77 further comprising the step:
2 requesting data associated with the data file utilizing the received pointer.
- 1 79. In a network including a client system and a plurality of server systems, a method for
2 enabling distributed program execution, the method comprising the steps:
3 (a) transmitting a graphical depiction of a data file; and
4 (b) providing a request to execute a first executable program for receiving application
5 output data on a client system, said request including information concerning a
6 second executable program for execution on one of the client system and one of a
7 plurality of server systems.
- 1 80. The method of claim 79 further comprising the step:
2 providing the first executable program.
- 1 81. The method of claim 79 further comprising the step:
2 providing a pointer to the data file.

- 62 -

- 1 82. The method of claim 81 further comprising the step:
2 servicing requests for data associated with the data file utilizing the provided pointer.
- 1 83. The method of claim 79 further comprising the step:
2 providing the contents of the data file.
- 1 84. The method of claim 79 further comprising the step:
2 providing a rule for determining whether the second executable program is to be executed
3 on the client system or one of the plurality of server systems.
- 1 85. The method of claim 79 further comprising the step:
2 providing a rule for determining whether the second executable program is to be executed
3 on the client system or one of the plurality of server systems that specifies the
4 determination between the client system and one of the plurality of server systems as
5 a function of at least one of the load on the client system, the load on one of the
6 plurality of server systems, the load on the plurality of server systems, and the speed
7 of the connection between the client system and the plurality of server systems.
- 1 86. In a network including a client system and a plurality of server systems, a method for
2 enabling distributed program execution, the method comprising the steps:
3 (a) transmitting, by a web system, a graphical depiction of a data file stored on a web
4 server;
5 (b) providing, to the client system, a request to execute a first executable program for
6 receiving application output data on a client system, said request including
7 information concerning a second executable program for execution on one of the
8 client system and one of a plurality of server systems; and
9 (c) executing, by the client system, the first executable program on the client system.
- 1 87. The method of claim 86 further comprising the step:
2 providing, to the client system, the first executable program.
- 1 88. The method of claim 86 further comprising the step:

- 63 -

2 receiving, at the client system, a rule for determining whether the second executable
3 program is to be executed on the client system or one of the plurality of server
4 systems.

1 89. The method of claim 86 further comprising the step:
2 receiving, at the client system, a rule for determining whether the second executable
3 program is to be executed on the client system or one of the plurality of server
4 systems that specifies the determination between the client system and one of the
5 plurality of server systems as a function of at least one of the load on the client
6 system, the load on one of the plurality of server systems, the load on the plurality of
7 server systems, and the speed of the connection between the client system and the
8 plurality of server systems.

1 90. The method of claim 86 further comprising the step:
2 (d) providing, to one of a plurality of server systems, a request to execute the second
3 executable program associated with the selected data file.

1 91. The method of claim 90 wherein step (d) comprises the steps:
2 (d-a) evaluating, by the client system, a rule to determine whether the second executable
3 program is to be executed on the client system or one of the plurality of server
4 systems; and
5 (d-b) sending, by the client system, a request to execute the second executable program
6 on one of the plurality of server systems based on the evaluation in step (d-a).

1 92. The method of claim 86 further comprising the step:
2 receiving, by the client system, application output data from the second executable
3 program.

1 93. The method of claim 92 further comprising the step:
2 displaying, by the client system, the application output data using the first executable
3 program.

1 94. The method of claim 86 further comprising the step:
2 providing, by the web server, data associated with the selected data file.

- 64 -

- 1 95. The method of claim 86 further comprising the step:
2 providing, by the web server, a pointer to the selected data file.
- 1 96. The method of claim 95 further comprising the step:
2 requesting, by the client system, data associated with the selected data file utilizing the
3 provided pointer.




FIG. 1

2/17




FIG. 2A




FIG. 2B

3/17






FIG. 2C

4/17



5/17




FIG. 3C

6/17




FIG. 3D

7/17




FIG. 4

8/17




FIG. 5

9/17




FIG. 6A

10/17


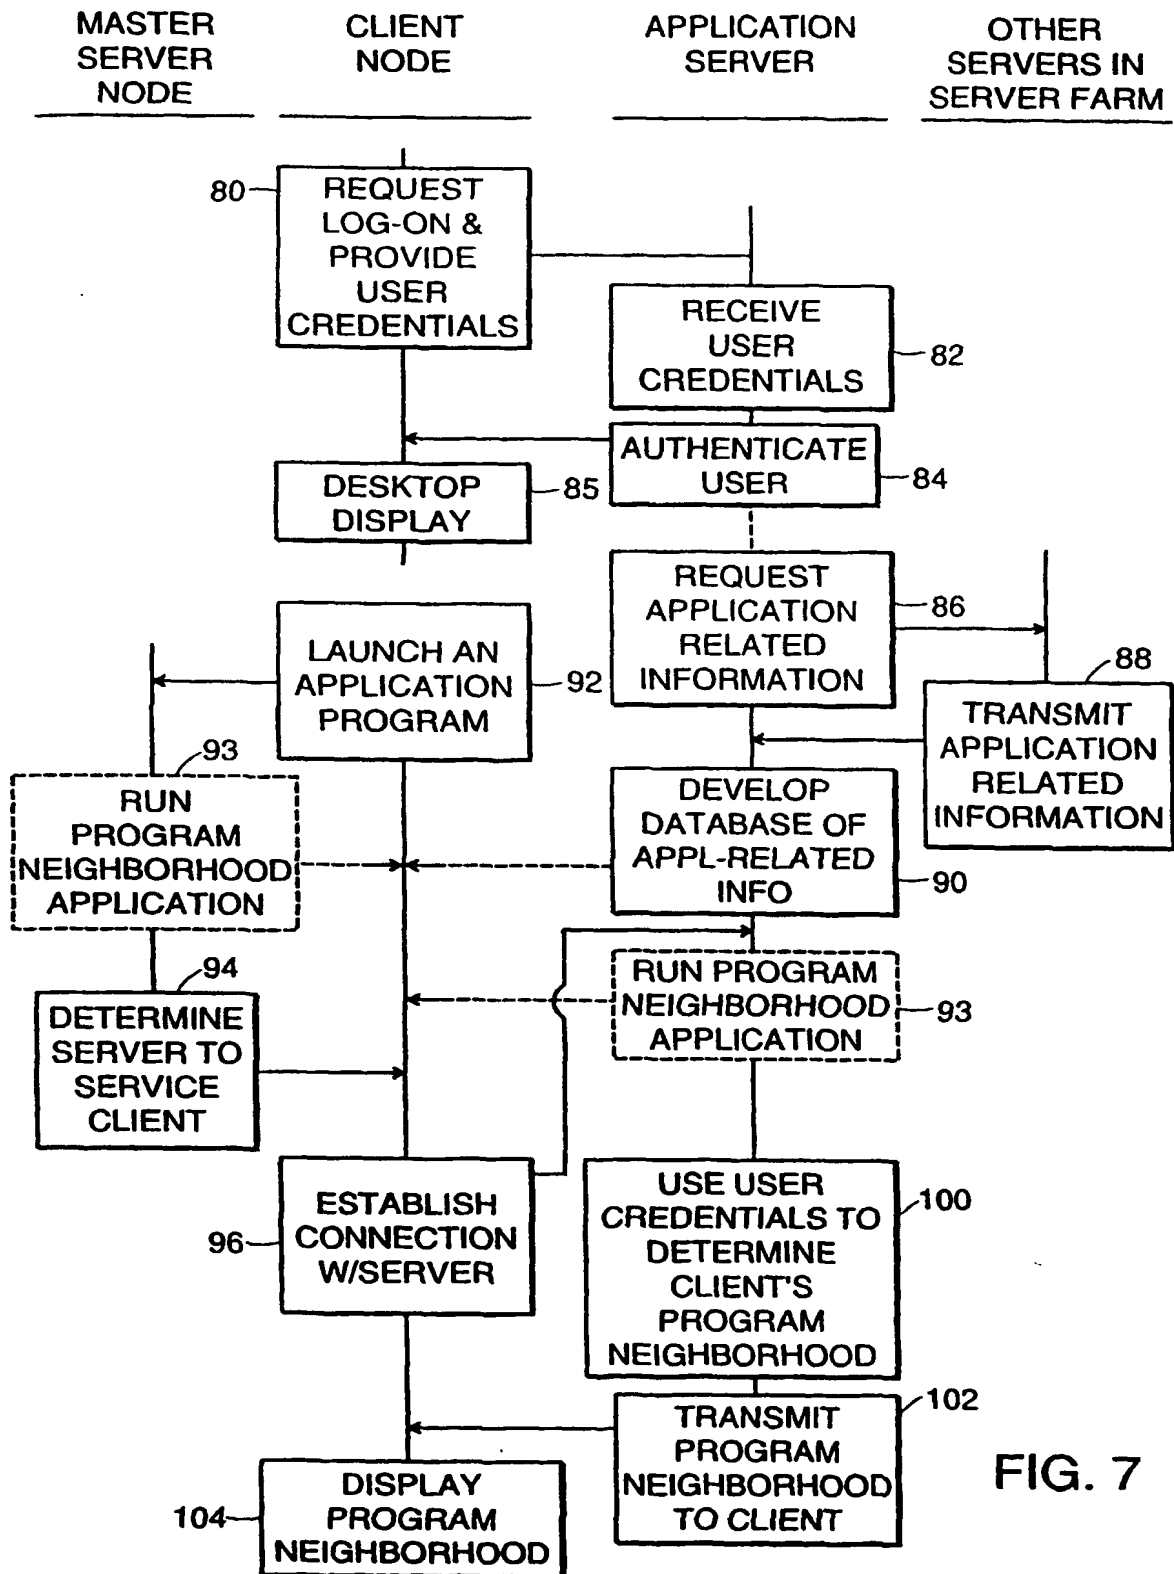


FIG. 6B

11/17



12/17




FIG. 8A

13/17

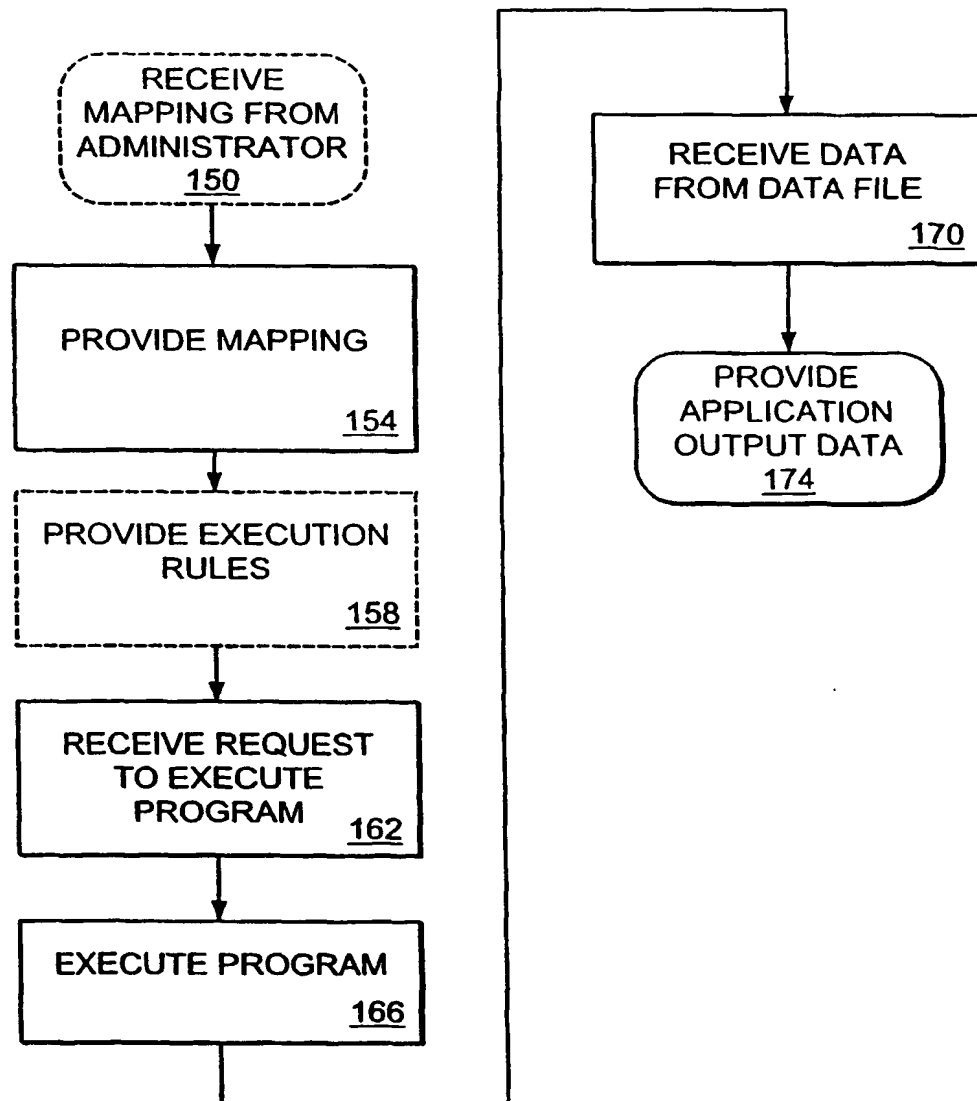


FIG. 8B

14/17




FIG. 9A

15/17

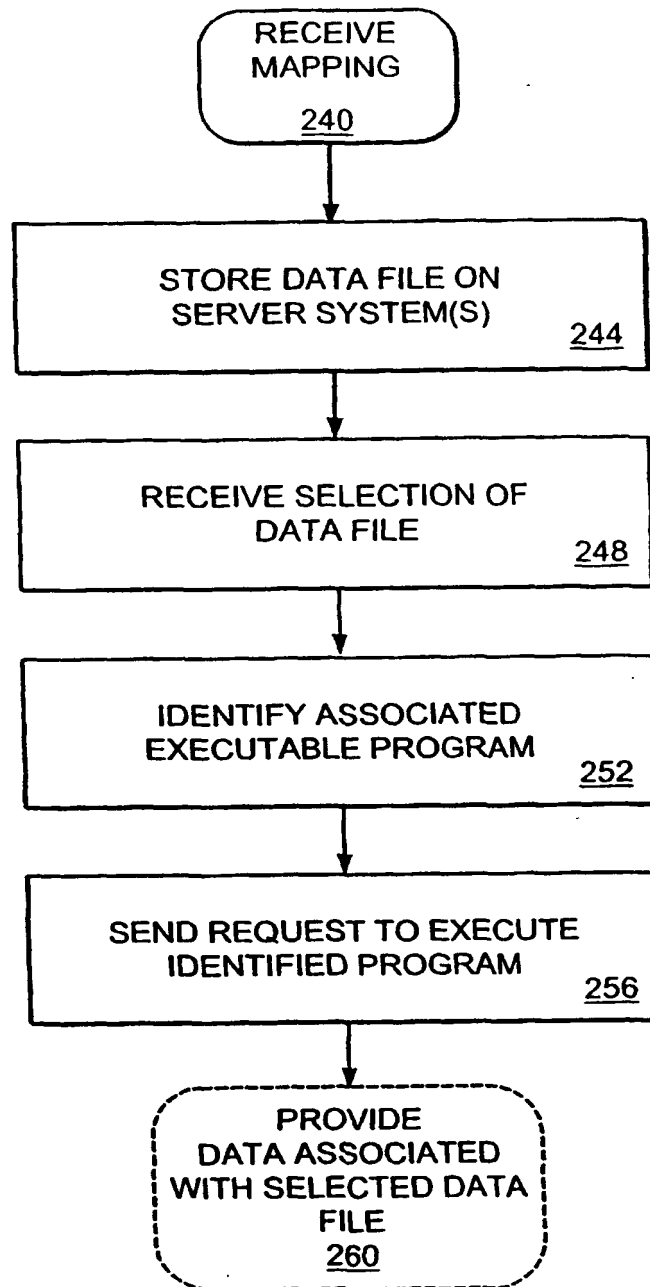


FIG. 9B

16/17




FIG. 10A

17/17

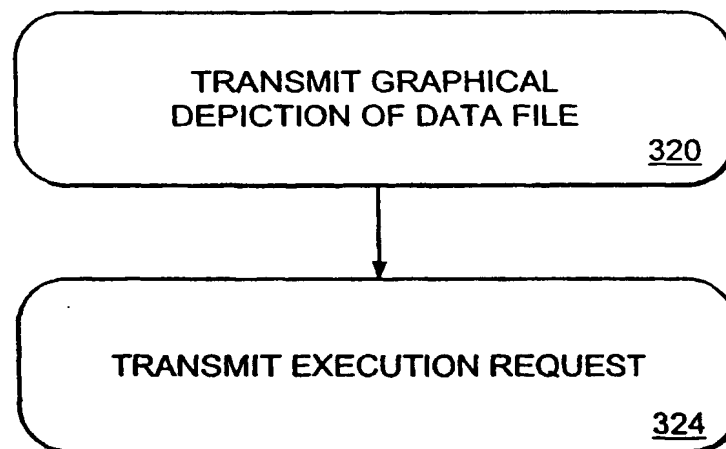


FIG. 10B

(19) World Intellectual Property
Organization
International Bureau



(43) International Publication Date
10 April 2003 (10.04.2003)

PCT

(10) International Publication Number
WO 2003/029977 A3

(51) International Patent Classification⁷: **G06F 9/50**

(21) International Application Number:
PCT/US2002/031261

(22) International Filing Date: 1 October 2002 (01.10.2002)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
09/970,038 2 October 2001 (02.10.2001) US
09/970,037 2 October 2001 (02.10.2001) US
09/970,462 2 October 2001 (02.10.2001) US

(71) Applicant: CITRIX SYSTEMS, INC. [US/US]; 851 W.
Cypress Creek Road, Fort Lauderdale, FL 33309 (US).

(72) Inventors: PEART, Franklyn; 3371 Overlook Road,
Davie, FL 33328 (US). RICHTBERG, Michael; 5461
NW 86th Terrace, Coral Springs, FL 33067 (US).

(74) Agent: LANZA, John, D.; Testa, Hurwitz & Thibault,
LLP, High Street Tower, 125 High Street, Boston, MA
02110 (US).

(81) Designated States (*national*): AE, AG, AL, AM, AT, AU,
AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU,
CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH,
GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC,
LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW,
MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG,
SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VC,
VN, YU, ZA, ZM, ZW.

(84) Designated States (*regional*): ARIPO patent (GH, GM,
KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW),
Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM),
European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE,
ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, SK,
TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ,
GW, ML, MR, NE, SN, TD, TG).

Published:

- with international search report
- before the expiration of the time limit for amending the
claims and to be republished in the event of receipt of
amendments

(88) Date of publication of the international search report:
19 August 2004

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: METHODS FOR DISTRIBUTED PROGRAM EXECUTION WITH FILE-TYPE ASSOCIATION IN A CLIENT-SERVER NETWORK

(57) Abstract: In one aspect, a mapping between a file type and a program is received, a client system file is graphically presented, a selection of the graphical depiction is received, a program associated with the selected file is identified, and a request to execute the identified program is sent. In another aspect, a server system file is graphically presented, a selection of the graphical depiction is received, the selection is transmitted to a server system, a request to execute an associated program is received, and the associated program is executed. In still another aspect, a web server file is graphically presented and a request to execute a first program that includes information concerning a second program associated with the file is received. Application output data from the second program is reviewed and formatted by the first program and received application output data is presented to the user.

WO 2003/029977 A3

INTERNATIONAL SEARCH REPORT

PCT/US 02/31261

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 G06F9/50

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 7 G06F H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, WPI Data, INSPEC, IBM-TDB

C. DOCUMENTS CONSIDERED TO BE RELEVANT

| Category * | Citation of document, with indication, where appropriate, of the relevant passages | Relevant to claim No. |
|------------|---|---|
| X | <p>US 5 838 906 A (ANG CHEONG S ET AL) 17 November 1998 (1998-11-17)</p> <p>the whole document</p> <p>---</p> <p>-/--</p> | <p>1-6, 11-21, 24-34, 39-48, 53-60, 65-68, 73-83, 86,87, 89,90, 92-96</p> |

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents :

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

- *T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- *X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- *Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- *Z* document member of the same patent family

Date of the actual completion of the international search

13 May 2004

Date of mailing of the international search report

06/07/2004

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Kielhöfer, P

INTERNATIONAL SEARCH REPORT

'US 02/31261

| C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT | | |
|--|--|--|
| Category * | Citation of document, with indication, where appropriate, of the relevant passages | Relevant to claim No. |
| X | <p>ANONYMOUS: "Tarantella Web-enabling Software, One World, One Network, One Answer, An SCO Technical White Paper " SANTA CRUZ OPERATION, INC., 'Online! 31 December 1999 (1999-12-31), pages 1-25, XP002251538 Retrieved from the Internet: <URL:http://web.archive.org/web/19990915181007/http://www.tarantella.sco.com/info/wps/wp1.pdf > 'retrieved on 2003-08-18!</p> <p>the whole document</p> | <p>1-6, 13-17, 24-29, 39, 46-48, 56,65, 67,68, 71,73, 74,79, 80,86, 87,90, 92-94</p> |
| A | <p>ANONYMOUS: "Tarantella Administration Guide > Reference > Objects and attributes : Windows Application object" THE SANTA CRUZ OPERATION, INC., 'Online! 1999, pages 1-4, XP002277254 Retrieved from the Internet: <URL:http://www.tarantella.com/knowhow/ei/help/html/en_US/ag_ref/windows_object.htm l> 'retrieved on 2004-04-19! the whole document</p> | <p>7-10,22, 23, 35-38, 49-52, 61-64, 69-72, 84,85, 88-91</p> |
| A | <p>US 6 128 642 A (DORASWAMY NAGANAND ET AL) 3 October 2000 (2000-10-03)</p> <p>abstract column 2, line 46 -column 3, line 45</p> | <p>7-11,18, 22,23, 31, 35-38, 49-53, 61-64, 69-72, 84,85, 88-91</p> |
| A | <p>ANONYMOUS: "Tarantella Administration Guide > Introduction > Essential information for Tarantella Administrators : Giving users access to applications and documents" THE SANTA CRUZ OPERATION, INC., 'Online! 1999, pages 1-3, XP002277255 Retrieved from the Internet: <URL:http://www.tarantella.com/knowhow/ei/help/html/en_US/ag_intro/access.html> 'retrieved on 2004-04-19! the whole document</p> | <p>1-3,6,7, 13,14, 17,22, 24-26, 28,35, 39, 46-49, 56,61, 67,69, 79,84, 86,88</p> |

-/--

INTERNATIONAL SEARCH REPORT

PCT/US 02/31261

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

| Category * | Citation of document, with indication, where appropriate, of the relevant passages | Relevant to claim No. |
|------------|--|--|
| A | <p>ANONYMOUS: "Tarantella Administration Guide > Reference > Load balancing with arrays : Multiple application servers" THE SANTA CRUZ OPERATION, INC. , 'Online! 1999, pages 1-2, XP002277256 Retrieved from the Internet: <URL:http://www.tarantella.com/knowhow/e11/help/html/en_US/ag_ref/load_hosts.html> 'retrieved on 2004-04-19!</p> <p>the whole document</p> | <p>1,7-10, 13, 22-24, 29, 35-39, 46, 49-52, 56, 61-64, 67, 69-72, 84,85, 88-91</p> |
| A | <p>WO 01/38973 A (GLAXO GROUP LTD ;SIMPKINS JOSEPH D (US); MURPHY STEPHANE (US); KEE) 31 May 2001 (2001-05-31) the whole document</p> | <p>1-96</p> |

INTERNATIONAL SEARCH REPORT

PCT/US 02/31261

| Patent document cited in search report | | Publication date | Patent family member(s) | Publication date |
|---|---|---------------------|----------------------------|---------------------|
| US 5838906 | A | 17-11-1998 | US 2003154261 A1 | 14-08-2003 |
| US 6128642 | A | 03-10-2000 | NONE | |
| WO 0138973 | A | 31-05-2001 | AT 249068 T | 15-09-2003 |
| | | | AU 1807801 A | 04-06-2001 |
| | | | CA 2396442 A1 | 31-05-2001 |
| | | | DE 60005051 D1 | 09-10-2003 |
| | | | EP 1234230 A2 | 28-08-2002 |
| | | | JP 2003515807 T | 07-05-2003 |
| | | | WO 0138973 A2 | 31-05-2001 |

This Page Blank (uspto)